arXiv:1704.02145v1 [cs.LO] 7 Apr 2017

A Fine-Grained Hierarchy of Hard Problems in the Separated Fragment

Marco Voigt

Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany, Saarbrücken Graduate School of Computer Science

Abstract

Recently, the separated fragment (SF) has been introduced and proved to be decidable. Its defining principle is that universally and existentially quantified variables may not occur together in atoms. The known upper bound on the time required to decide SF's satisfiability problem is formulated in terms of quantifier alternations: Given an SF sentence $\exists\vec{\mathbf{z}}\forall\vec{\mathbf{x}}_1\exists\vec{\mathbf{y}}_1\ldots\forall\vec{\mathbf{x}}_n\exists\vec{\mathbf{y}}_n.\,\psi$ in which ψ is quantifier free, satisfiability can be decided in nondeterministic n-fold exponential time. In the present paper, we conduct a more fine-grained analysis of the complexity of SF-satisfiability. We derive an upper and a lower bound in terms of the *degree ∂ of interaction of existential variables (short: degree)*—a novel measure of how many separate existential quantifier blocks in a sentence are connected via joint occurrences of variables in atoms. Our main result is the k-NExpTime-completeness of the satisfiability problem for the set $\mathrm{SF}_{\partial\leq k}$ of all SF sentences that have degree k or smaller. Consequently, we show that SF-satisfiability is non-elementary in general, since SF is defined without restrictions on the degree. Beyond trivial lower bounds, nothing has been known about the hardness of SF-satisfiability so far.

1 Introduction

In [17] the separated fragment (SF) of first-order logic with equality is introduced. Its defining principle is that universally and existentially quantified variables may not occur together in atoms. (Topmost existential quantifier blocks are exempt from this rule.) SF properly generalizes both the Bernays–Schönfinkel–Ramsey (BSR) fragment ($\exists^*\forall^*$-sentences with equality) and the relational monadic fragment without equality (MFO). Still, the satisfiability problem for SF is decidable.

In computational logic formulas are often classified based on the shape of quantifier prefixes. There is a wealth of results that separate decidable first-order formulas from undecidable ones in this fashion, see [3] for references. The definition of the BSR fragment is only one example. In the context of computational complexity, hierarchies are defined, such as the polynomial hierarchy, where the hardness of problems is assumed to grow with the number of quantifier alternations that are allowed to occur.

Although the definition of SF breaks with the paradigm of restricting quantifier prefixes, the known upper bound on the complexity of SF-satisfiability is again based on quantifier prefixes: Deciding whether an SF sentence $\varphi := \exists\vec{\mathbf{z}}\forall\vec{\mathbf{x}}_1\exists\vec{\mathbf{y}}_1\ldots\forall\vec{\mathbf{x}}_n\exists\vec{\mathbf{y}}_n.\psi$ with quantifier-free ψ is satisfiable requires a nondeterministic computing time that is at most n-fold exponential in the length of φ (cf. Theorem 17 in [17]). On the one hand, we complement this result with a corresponding lower bound in the present paper. That is, we show that SF-satisfiability is indeed non-elementary. On the other hand, we derive a refined upper bound that is based on the *degree ∂ of interaction of existential variables*. An overview of the resulting hierarchy of complete problems is depicted in Figure 1. Intuitively, φ exhibits a degree $\partial_\varphi = k$, if variables from k distinct existential quantifier blocks interact. We say that two variables x, y interact, if they occur together in at least one atom or if there is a third variable z that interacts with both x and y (i.e. the property is transitive). For instance, in

the SF sentence $\forall x_1 \exists y_1 v_1 \forall x_2 \exists y_2 v_2 \forall x_3 \exists y_3 v_3.\ \big(P(x_1, x_2, x_3) \wedge \neg Q(y_1, y_3)\big) \vee P(y_2, v_2, v_3) \vee \neg Q(y_3, v_1)$ the sets $\{y_1, y_3, v_1\}$ and $\{y_2, v_2, v_3\}$ form the maximal sets of interacting existential variables. Since each of these sets contains variables from at most two distinct quantifier blocks, the formula exhibits a degree $\partial = 2$.

In Section 4.1, and in particular in Theorem 13, we observe that the satisfiability problem for $\mathrm{SF}_{\partial \leq k}$—the set of all SF sentences φ with $\partial_\varphi \leq k$—lies in k-NExpTime. It is worth mentioning that this result adequately accounts for the known complexity of MFO-satisfiability. For every MFO sentence φ we trivially have $\partial_\varphi = 1$, since all occurring predicate symbols have an arity of at most one. Theorem 13 entails that MFO-satisfiability is in NExpTime, which is well known. Still, this bound is not reproducible with the analysis of the complexity of SF-satisfiability conducted in [17]. Apparently, non-elementary satisfiability problems are not very widespread among the

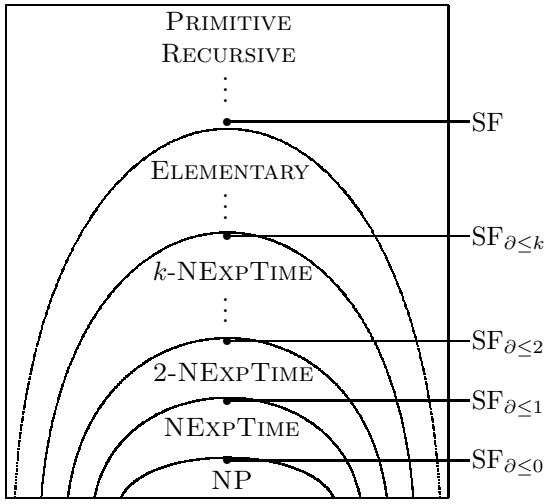

Figure 1: The subfragments of SF scale over the major nondeterministic complexity classes in Elementary, while SF itself goes even beyond.

decidable fragments of classical fist-order logic known today. We show in Section 5 that SF falls into this category. To the present author's knowledge, the only known companion in this respect is the fluted fragment (FL). Indeed, Pratt-Hartmann, Szwast, and Tendera show in [14] that satisfiability of fluted sentences with at most $2k$ variables is k-NExpTime-hard. Moreover, they argue that satisfiability of fluted sentences with at most k variables lies in k-NExpTime. Although a significant gap between these lower and upper bounds remains to be closed, the fluted fragment seems to comprise a similar hierarchy of hard problems as SF does.

Another characteristic of SF is that it enjoys a *small model property*. More precisely, given an SF sentence φ, one can compute a positive integer n that depends on the degree ∂_φ and the length of φ such that, if there is a model of φ at all, then there also is a model whose domain contains at most n elements. Many first-order fragments are known to enjoy a small model property. The BSR fragment and MFO are among the classical ones (see [3] for references). More recently defined fragments include the two-variable fragment (FO_2) [12], [8], the fluted fragment (FL) [15], [16], [14], the guarded fragment (GF) [1], [7], the guarded negation fragment (GNF) [2], and the uniform one-dimensional fragment (UF_1) [11]. While GNF and UF_1 are incomparable, GNF extends GF, and UF_1 can be considered as a generalization of FO_2. Guarded fragments and the two-variable fragment have received quite some attention due to the fact that modal logics have natural translations into them. As a continuation of that theme, we shall see in Section 3.1 how classes of sentences enjoying a small model property can be effectively translated into subclasses of SF. During the translation process the length of formulas increases by a factor that is logarithmic in the size of small models of the original. One benefit of translating non-SF sentences into SF

sentences is that in SF one can natively express concepts such as transitivity and basic counting quantifiers (Proposition 5). This is not always possible in other fragments enjoying a small model property. For example, transitivity cannot be expressed in FO_2, GF, and FL.

Summing up, the main contributions are:
(i) Based on the novel concept of the degree of interaction of existential variables, we substantially refine the existing analysis of the time required to decide SF-satisfiability. More concretely, we show that a satisfiable SF sentence φ with $\partial_\varphi = k$ has a model whose domain is of a size that is at most k-fold exponential in the length of φ (Section 4, Theorem 13). With this refined approach we can close the complexity gap for the class of *strongly separated* sentences (Corollary 15) that was left open in [17]. Moreover, the complexity of MFO can be explained in the refined framework.
(ii) We complement the complexity analysis with corresponding lower bounds in two respects. We first derive a lower bound on the length of shortest BSR sentences that are equivalent to a given SF sentence (Section 4.2, Theorem 16). In Section 5, we prove k-NExpTime-hardness of satisfiability for the class of SF sentences φ with $\partial_\varphi = k$ (Theorem 21). Since SF is in general defined without restrictions on the degree ∂_φ, our result implies that SF-satisfiability is non-elementary.
(iii) We devise a simple translation from classes of first-order sentences that enjoy a small model property into SF (Proposition 4). Moreover, we argue that SF can express basic counting quantifiers (Proposition 5).

In order to facilitate smooth reading, most proofs are only sketched in the main text and presented in full in the appendix. The present paper is the full version of the extended abstract [18].

2 Preliminaries

We mainly reuse the basic notions from [17]. We repeat the definition of necessary concepts and notation for the sake of completeness.

We consider first-order logic formulas. The underlying signature shall not be mentioned explicitly, but will become clear from the current context. For the distinguished *equality* predicate we use $\approx$. We follow the convention that negation binds strongest, that conjunction binds stronger than disjunction, and that all of the aforementioned bind stronger than implication. The scope of quantifiers shall stretch as far to the right as possible. By $\mathrm{len}(\cdot)$ we denote a reasonable measure of the length of formulas satisfying $\mathrm{len}(\varphi \to \psi) = \mathrm{len}(\neg\varphi \vee \psi)$ and $\mathrm{len}(\varphi \leftrightarrow \psi) = \mathrm{len}((\neg\varphi \vee \psi) \wedge (\varphi \vee \neg\psi))$.

We write $\varphi(x_1, \ldots, x_m)$ to denote a formula φ whose free variables form a subset of $\{x_1, \ldots, x_m\}$. In all formulas we tacitly assume that no variable occurs freely and bound at the same time and that no variable is bound by two different occurrences of quantifiers, unless explicitly stated otherwise. For convenience, we sometimes identify tuples $\vec{\mathbf{x}}$ of variables with the set containing all the variables that occur in $\vec{\mathbf{x}}$. We write $\mathrm{vars}(\varphi)$ to address the set of all variable symbols that occur in φ. Similarly, $\mathrm{consts}(\varphi)$ denotes the set of all constant symbols in φ. We denote *substitution* by $\varphi\big[x/t\big]$ if every free occurrence of x in φ is to be substituted with the term t.

A *literal* is an atom or a negated atom, and a *clause* is a disjunction of literals. We say that a formula is in *conjunctive normal form (CNF)*, if it is a conjunction of clauses, possibly preceded by a quantifier prefix. A formula in CNF is *Horn* if every clause contains at most one non-negated literal. It is *Krom* if every clause contains at most two literals at all.

A sentence $\varphi := \forall\vec{\mathbf{x}}_1 \exists\vec{\mathbf{y}}_1 \ldots \forall\vec{\mathbf{x}}_n \exists\vec{\mathbf{y}}_n.\psi$ is in *standard form*, if it is in *negation normal form* (i.e. every negation symbol occurs directly before an atom) and ψ is quantifier free, contains exclusively the Boolean connectives $\wedge, \vee, \neg$, and does not contain non-constant function symbols. The tuples $\vec{\mathbf{x}}_1$ and $\vec{\mathbf{y}}_n$ may be empty, i.e. the quantifier prefix does not have to start with a universal quantifier, and it does not have to end with an existential quantifier. Moreover, we require that every variable occurring in the quantifier prefix does also occur in ψ.

As usual, we interpret a formula φ with respect to given structures. A *structure* $\mathcal{A}$ consists of a nonempty *universe* $\mathfrak{U}_\mathcal{A}$ (also: *domain*) and interpretations $f^\mathcal{A}$ and $P^\mathcal{A}$ of all considered function and predicate symbols, respectively, in the usual way. Given a formula φ, a structure $\mathcal{A}$, and a variable assignment β, we write $\mathcal{A}, \beta \models \varphi$ if φ evaluates to *true* under $\mathcal{A}$ and β. We write $\mathcal{A} \models \varphi$ if $\mathcal{A}, \beta \models \varphi$ holds for every β. The symbol $\models\!\!\models$ denotes *(semantic) equivalence* of formulas, i.e. $\varphi \models\!\!\models \psi$

holds whenever for every structure $\mathcal{A}$ and every variable assignment β we have $\mathcal{A}, \beta \models \varphi$ if and only if $\mathcal{A}, \beta \models \psi$. We call two sentences φ and ψ *equisatisfiable* if φ has a model if and only if ψ has one.

A structure $\mathcal{A}$ is a *substructure* of a structure $\mathcal{B}$ (over the same signature) if (1) $\mathfrak{U}_\mathcal{A} \subseteq \mathfrak{U}_\mathcal{B}$, (2) $c^\mathcal{A} = c^\mathcal{B}$ for every constant symbol c, (3) $P^\mathcal{A} = P^\mathcal{B} \cap \mathfrak{U}_\mathcal{A}^m$ for every m-ary predicate symbol P, and (4) $f^\mathcal{A}(\mathfrak{a}_1, \ldots, \mathfrak{a}_m) = f^\mathcal{B}(\mathfrak{a}_1, \ldots, \mathfrak{a}_m)$ for every m-ary function symbol f and every m-tuple $\langle \mathfrak{a}_1, \ldots, \mathfrak{a}_m \rangle \in \mathfrak{U}_\mathcal{A}^m$. The following is a standard lemma, see, e.g., [5] for a proof.

Lemma 1 (Substructure lemma)**.** Let φ be a first-order sentence in prenex normal form without existential quantifiers and let $\mathcal{A}$ be a substructure of $\mathcal{B}$. $\mathcal{B} \models \varphi$ entails $\mathcal{A} \models \varphi$.

Lemma 2 (Miniscoping)**.** Let φ, ψ, χ be arbitrary first-order formulas, and assume that x does not occur freely in χ.

$$\begin{array}{lcl} \exists x.(\varphi \vee \psi) & \models\!\!\!\models & (\exists x_1.\varphi) \vee (\exists x_2.\psi) \ , \\ \exists x.(\varphi \circ \chi) & \models\!\!\!\models & (\exists x.\varphi) \circ \chi \ , \\ \forall x.(\varphi \wedge \psi) & \models\!\!\!\models & (\forall x_1.\varphi) \wedge (\forall x_2.\psi) \ , \\ \forall x.(\varphi \circ \chi) & \models\!\!\!\models & (\forall x.\varphi) \circ \chi \ , \end{array}$$

where $\circ \in \{\wedge, \vee\}$.

We use the notation $[k]$ to abbreviate the set $\{1, \ldots, k\}$ for any positive integer k. Moreover, $\mathfrak{P}$ shall be used as the power set operator, i.e. $\mathfrak{P}S$ denotes the set of all subsets of a given set S. Finally, we need some notation for the *tetration operation*. We define $2^{\uparrow k}(m)$ inductively: $2^{\uparrow 0}(m) := m$ and $2^{\uparrow k+1}(m) := 2^{(2^{\uparrow k}(m))}$.

3 The separated fragment

Let φ be a first-order formula. We call two disjoint sets of variables X and Y *separated in* φ if and only if for every atom A occurring in φ we have $\mathrm{vars}(A) \cap X = \emptyset$ or $\mathrm{vars}(A) \cap Y = \emptyset$.

Definition 3 (Separated fragment (SF), [17])**.** The *separated fragment (SF)* of first-order logic consists of all existential closures of prenex formulas without non-constant function symbols in which existentially quantified variables are separated from universally quantified ones. More precisely, SF consists of all first-order sentences with equality but without non-constant function symbols that are of the form $\exists \vec{\mathbf{z}} \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n . \psi$, in which ψ is quantifier-free, and in which the sets $\vec{\mathbf{x}}_1 \cup \ldots \cup \vec{\mathbf{x}}_n$ and $\vec{\mathbf{y}}_1 \cup \ldots \cup \vec{\mathbf{y}}_n$ are separated.

The tuples $\vec{\mathbf{z}}$ and $\vec{\mathbf{y}}_n$ may be empty, i.e. the quantifier prefix does not have to start with an existential quantifier and it does not have to end with an existential quantifier either.

Notice that the variables in $\vec{\mathbf{z}}$ are not subject to any restriction concerning their occurrences.

In [17] the authors show that the satisfiability problem for SF sentences (*SF-satisfiability*) is decidable. Before we start investigating the complexity issues related to SF-satisfiability, we elaborate on the expressiveness of SF.

3.1 Expressiveness

Every SF sentence is equivalent to a BSR sentence ([17], Lemma 6). We shall outline in Section 4 how to analyze the blow-up that we have to incur during this translation process and how it depends on the degree of interaction of existential variables. Since the BSR fragment enjoys a small model property (cf. Proposition 6), SF inherits the small model property from BSR. However, regarding the size of minimal models of satisfiable formulas, SF sentences are much more compact. While satisfiable BSR sentences have models whose domain is linear in the length of the formula, satisfiable SF sentences can enforce domains of a size that cannot be bounded from above by a finite tower of exponentials. We provide first evidence for this fact in Theorem 16, where we give a non-elementary lower bound on the length of equivalent BSR sentences. This lower bound even applies to the SF-Horn-Krom subfragment of SF. Moreover, we exploit the capability of SF

sentences φ to enforce models of ∂_φ-fold exponential size in the proof of the k-NExpTime-hardness of SF-satisfiability (for every $k \geq 1$).

Apart from compactness of representation, and from the perspective of satisfiability, all first-order fragments that enjoy small model properties share a common ground of expressiveness. Neglecting efficiency, every sentence φ from such a fragment can be effectively translated into a (finite) propositional formula ϕ in such a way that from a satisfying variable assignment for ϕ one can straightforwardly reconstruct a (Herbrand) model of φ. The reason is simply that universal quantification can then be understood as finite conjunction (over a finite domain) and existential quantification can be conceived as finite disjunction.

The following proposition illustrates why SF is to some extent prototypical for first-order fragments that enjoy a small model property.

Proposition 4. Consider any nonempty class $\mathcal{C}$ of first-order formulas without non-constant function symbols for which we know a computable mapping $f : \mathcal{C} \to \mathbb{N}$ such that every satisfiable φ in $\mathcal{C}$ has a model of size at most $f(\varphi)$. Then there exists an effective translation T from $\mathcal{C}$ into SF such that for every $\varphi \in \mathcal{C}$ (a) every model of $T(\varphi)$ is also a model of φ, (b) every model of φ whose size is at most $f(\varphi)$ can be extended to a model of $T(\varphi)$ over the same domain, and (c) the length of $T(\varphi)$ lies in $\mathcal{O}\big(\mathrm{len}(\varphi) \cdot \log f(\varphi) \cdot \log\log f(\varphi)\big)$.

Proof. We outline the translation T for some given input sentence φ, which we assume to be in negation normal form (without loss of generality). Let $m := \lceil \log_2 f(\varphi) \rceil$ and let $Q_1, \ldots, Q_m$ be unary predicate symbols that do not occur in φ. For all terms s, t we define $s \mathrel{\hat{\approx}} t$ as abbreviation of $\bigwedge_{i=1}^{m} Q_i(s) \leftrightarrow Q_i(t)$. In order to restrict the domain to 2^m elements, we conjoin the formula $\chi_{\mathrm{fin}} := \forall xy.\, x \mathrel{\hat{\approx}} y \to x \approx y$. Since in any structure $\mathcal{A}$ there are at most 2^m domain elements that can be distinguished by their membership in the sets $Q_1^{\mathcal{A}}, \ldots, Q_m^{\mathcal{A}}$, it is clear that $\mathcal{A} \models \chi_{\mathrm{fin}}$ entails $|\mathfrak{U}_{\mathcal{A}}| \leq 2^m$. Moreover, we observe the following property.

(∗) Let $\mathcal{A}$ be any structure, let β be any variable assignment over $\mathcal{A}$'s domain, and let s, t be two terms. If $\mathcal{A} \models \chi_{\mathrm{fin}}$ holds, then we get $\mathcal{A}, \beta \models s \mathrel{\hat{\approx}} t$ if and only if $\mathcal{A}, \beta \models s \approx t$.

This means, if we restrict our attention to domains with at most 2^m domain elements, we can now use a separated form of equality.

(∗∗) Let ψ be any first-order formula and let v be some variable that does not occur in ψ. Then ψ is equivalent to $\forall v.\, u \approx v \to \psi\big[u/v\big]$.

We can transform φ into an equivalent sentence φ' by consecutively replacing each subformula of the form $\exists y.\psi$ in φ with $\exists y \forall v.\, y \approx v \to \psi\big[y/v\big]$, where we assume v to be fresh (one fresh variable for every replaced subformula). Consequently, every atom in φ' that is not an equation contains exclusively universally quantified variables. Moreover, (∗∗) implies that φ and φ' are equivalent.

Let φ'' be the result of replacing all equations $y \approx v, v \approx y$ in φ' in which y is existentially quantified and v universally quantified with the formula $y \mathrel{\hat{\approx}} v$. We then set $\varphi_{\mathrm{SF}} := \chi_{\mathrm{fin}} \wedge \varphi''$. By (∗), any model of φ_{SF} is also a model of φ. Conversely, any model $\mathcal{A}$ of φ that has at most 2^m domain elements can be converted into a model $\mathcal{B}$ of φ_{SF} by defining the relations $Q_1^{\mathcal{B}}, \ldots, Q_m^{\mathcal{B}}$ in an appropriate way. □

The φ_{SF} in the above proof belongs to a subfragment of SF that we call *strongly separated* (cf. Definition 14) and whose satisfiability problem is complete for NExpTime (cf. Corollary 15).

Unfortunately, the translation methodology of Proposition 4 does not help in the quest for new decidable first-order fragments. The reason is simply that we already need arguments leading to a small model property before we can start the translation process, as we need information about the size of the models that have to be considered. Nevertheless, such translations can be useful in view of the expressiveness of SF that other first-order fragments, such as FO_2, the fluted fragment, and GF, lack. For instance, SF sentences can naturally express the axioms of equivalence, most prominently, transitivity. Hence, fundamental and interesting properties of predicates that have to be assumed at the meta-level when dealing with less expressive logics can be formalized directly

in SF. Moreover, basic counting quantifiers can be defined natively in SF and do not have to be introduced via special operators. More precisely, given any formula $\exists^{\geq n} y.\, \varphi$ with positive n and without non-constant function symbols, its standard translation $\exists y_1 \dots y_n.\, \bigwedge_{i=1}^{n} \varphi\big[y/y_i\big] \wedge \bigwedge_{i<j} y_i \not\approx y_j$ is not in conflict with the separateness conditions of SF's definition, if the set $\{y_1, \dots, y_n\}$ is separated in φ from the set of universally quantified variables.

Proposition 5. Counting quantifiers $\exists^{\geq n}$ with positive integer n are expressible in SF.

3.2 Basic complexity considerations

We first recall the well-known small model properties of SF's subfragments BSR and MFO (see [3] for references).

Proposition 6. Let $\varphi := \exists \vec{\mathbf{z}} \forall \vec{\mathbf{x}}.\psi$ be a satisfiable BSR sentence, i.e. ψ is quantifier free and does not contain non-constant function symbols. There is a model $\mathcal{A} \models \varphi$ such that $|\mathfrak{U}_{\mathcal{A}}| \leq \max\big(|\vec{\mathbf{z}}| + |\mathrm{consts}(\varphi)|, 1\big)$.

We make use of this property when we derive an upper bound on the size of small models for satisfiable SF sentences, as our approach will be based on an effective translation of SF sentences into equivalent BSR sentences.

Proposition 7. Let $\varphi := \exists \vec{\mathbf{z}} \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \dots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n.\psi$ be a satisfiable monadic sentence without equality and without non-constant function symbols, i.e. all predicate symbols in φ are of arity 1. Moreover, assume that φ contains k distinct predicate symbols. There is a model $\mathcal{A} \models \varphi$ such that $|\mathfrak{U}_{\mathcal{A}}| \leq 2^k$.

Notice that the shape of the quantifier prefix does not contribute to the upper bound.

The following lemma links bounds on the size of models with the computing time that is required to decide satisfiability.

Lemma 8 (cf. [3], Proposition 6.0.4)**.** Let φ be a first-order sentence in prenex normal form containing n universally quantified variables. The question whether φ has a model of cardinality m can be decided nondeterministically in time $p\big(m^n \cdot \mathrm{len}(\varphi)\big)$ for some polynomial p.

With this lemma at hand, it is enough to prove a small model property for a given class of first-order sentences, in order to bound the worst-case time complexity of the corresponding satisfiability problem from above. This is exactly what the authors of [17] have done for SF.

Proposition 9 ([17], Theorem 17)**.** Let $\varphi := \exists \vec{\mathbf{z}} \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \dots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n.\psi$ be an SF sentence for some quantifier-free ψ. There is some equivalent BSR sentence $\exists \vec{\mathbf{u}} \forall \vec{\mathbf{v}}.\psi'$ in which the number of occurring constant symbols plus the number of existential quantifiers is at most $\mathrm{len}(\varphi) + n \cdot \mathrm{len}(\varphi) \cdot \big(2^{\uparrow n}(\mathrm{len}(\varphi))\big)^n$. As a result, satisfiability of φ can be decided nondeterministically in time that is at most n-fold exponential in $\mathrm{len}(\varphi)$.

Clearly, applying this result to an MFO sentence substantially overshoots the actual worst-case time requirements. To stress it again, the notion of the degree of interaction is a remedy to this sort of inaccuracies, as we shall see in Section 4.

A special case that is worth considering, before we investigate the complexity of full SF, is the class of SF sentences that do not contain universal quantifiers. This species of formulas coincides with the purely existential fragment of first-order logic without non-constant function symbols, and it is a close relative of propositional logic. Recall that SAT is NP-complete [4], Horn-SAT is P-complete [10, 13], and 2SAT is NL-complete [9].

Proposition 10.

(i) Satisfiability for the class of SF sentences without universal quantifiers is NP-complete.

(ii) Satisfiability for the class of SF-Horn sentences without universal quantifiers is P-complete.

(iii) Satisfiability for the class of SF-Krom sentences without universal quantifiers and without equality is NL-complete.

Proof sketch. The proof of (i) – (iii) proceeds by reductions to the corresponding satisfiability problems for propositional logic and back. This is straightforwardly done by Skolemization as long as we consider only SF sentences without equality (cf. Lemma 26 in the appendix).

In the latter cases, we first Skolemize exhaustively, producing φ_{gnd}, which is ground and contains only Skolem constants and no non-constant function symbols. Then we use the standard trick to eliminate the equality predicate $\approx$. We introduce a fresh binary relation symbols E and replace every equation $c \approx d$ with an atom $E(c, d)$. Moreover, we add axioms for reflexivity, symmetry, transitivity, and congruence. Of course, we do not use the universally quantified axioms but rather add their instances with respect to all the constant symbols that occur in φ_{gnd}. To avoid an exponential blow-up in the case of the congruence axioms, we only add the instances that affect non-equational atoms which really occur in φ_{gnd}.

Let ϕ be the propositional formula that results from φ_{gnd} by replacing every ground atom A with the propositional variable p_A. We observe that $\mathrm{len}(\phi) \in \mathcal{O}\big(\mathrm{len}(\varphi)^3\big)$. Moreover, if φ is a Horn formula, then ϕ is Horn. Notice that the outlined elimination of equality does not preserve the Krom property. □

4 Translation of SF sentences into BSR sentences

In this section, we analyze the transformation process from SF into the BSR fragment from the perspective of the *degree of interaction of existential variables*. Our aim is to derive upper and lower bounds on the length of the resulting BSR-formulas. Roughly speaking, in the first phase of the translation process all quantifiers are moved inwards as far as possible (cf. the proof of Lemma 12). In order to do so, we first transform the sentence in question into a formula in CNF. After that, we employ the well-known rules of miniscoping (cf. Lemma 2), supplemented by the rule formulated in the following lemma.

Lemma 11. Let I and K_i, $i \in I$, be sets that are finite, nonempty, and pairwise disjoint. The elements of these sets serve as indices. Let

$$\varphi := \exists \vec{\mathbf{y}}. \bigwedge_{i \in I} \Big(\chi_i(\vec{\mathbf{z}}) \vee \bigvee_{k \in K_i} \eta_k(\vec{\mathbf{y}}, \vec{\mathbf{z}}) \Big)$$

be some first-order formula where the χ_i and the η_k denote arbitrary subformulas that we treat as indivisible units in what follows. We say that $f : I \to \bigcup_{i \in I} K_i$ is a *selection function* if for every $i \in I$ we have $f(i) \in K_i$. We denote the set of all selection functions of this form by $\mathcal{F}$.

Then φ is equivalent to $\varphi' :=$

$$\bigwedge_{\substack{S \subseteq I \\ S \neq \emptyset}} \Big(\bigvee_{i \in S} \chi_i(\vec{\mathbf{z}}) \Big) \vee \bigvee_{f \in \mathcal{F}} \Big(\exists \vec{\mathbf{y}}. \bigwedge_{i \in S} \eta_{f(i)}(\vec{\mathbf{y}}, \vec{\mathbf{z}}) \Big) .$$

Proof sketch. The proof of this lemma follows a conceptually simple strategy. Using the distributivity of $\wedge$ over $\vee$, we first transform φ into a disjunction of conjunctions of the indivisible units $\chi_i(\vec{\mathbf{z}})$ and $\eta_k(\vec{\mathbf{y}}, \vec{\mathbf{z}})$. Then, exploiting the equivalences in Lemma 2, we push the existential quantifier block $\exists \vec{\mathbf{y}}$ inwards such that it only binds conjunctions of units $\eta_k(\vec{\mathbf{y}}, \vec{\mathbf{z}})$. This is possible, because none of the variables in $\vec{\mathbf{y}}$ occurs in any of the $\chi_i(\vec{\mathbf{z}})$. From this point on, we treat the newly emerged subformulas $\exists \vec{\mathbf{y}}. \bigwedge_{k'} \eta_{k'}(\vec{\mathbf{y}}, \vec{\mathbf{z}})$ as if they were indivisible. We then transform the formula back into a conjunction of disjunctions of indivisible units, this time using the distributivity of $\vee$ over $\wedge$. It then remains to show that the result of this transformation exhibits a highly redundant structure and is actually equivalent to φ'. □

4.1 Degree of interaction of existential variables and the size of small models

Consider the formula $\varphi := \exists\vec{\mathbf{z}}\forall\vec{\mathbf{x}}_1\exists\vec{\mathbf{y}}_1 \ldots \forall\vec{\mathbf{x}}_n\exists\vec{\mathbf{y}}_n.\psi$ in standard form in which ψ is quantifier free and in which the sets $\vec{\mathbf{x}} := \vec{\mathbf{x}}_1 \cup \ldots \cup \vec{\mathbf{x}}_n$ and $\vec{\mathbf{y}} := \vec{\mathbf{y}}_1 \cup \ldots \cup \vec{\mathbf{y}}_n$ are separated. In addition, we assume that $\vec{\mathbf{x}}_1$ and $\vec{\mathbf{y}}_1$ are nonempty. The tuple $\vec{\mathbf{z}}$, on the other hand, may be empty.

For any $j \in [n]$ and any variable $y \in \vec{\mathbf{y}}_j$ we say that y *is a level-j variable*, denoted $\mathrm{lvl}(y) = j$. For any nonempty set $Y \subseteq \vec{\mathbf{y}}$ of existentially quantified variables and any positive integer k we say that Y *has degree k in φ*, denoted $\partial_{Y,\varphi} = k$, if k is the maximal number of distinct variables $y_1, \ldots, y_k \in Y$ that belong to different levels in φ, i.e. $\mathrm{lvl}(y_1) < \ldots < \mathrm{lvl}(y_k)$. We say that φ*'s degree of interaction of existential variables* (short: *degree*) is k, denoted $\partial_\varphi = k$, if k is the smallest positive integer such that we can partition $\vec{\mathbf{y}}$ into $m > 0$ parts $Y_1, \ldots, Y_m$ that are all pairwise separated in φ and for which $k = \max\{k_j \mid \partial_{Y_j,\varphi} = k_j, 1 \leq j \leq m\}$. Sentences $\varphi := \exists\vec{\mathbf{z}}\forall\vec{\mathbf{x}}.\psi$ in standard form with quantifier-free ψ are said to have *degree* zero, i.e. $\partial_\varphi = 0$, if $\vec{\mathbf{x}}$ is empty and we define $\partial_\varphi = 1$ if $\vec{\mathbf{x}}$ is nonempty.

Lemma 12. Let $\varphi := \exists\vec{\mathbf{z}}\forall\vec{\mathbf{x}}_1\exists\vec{\mathbf{y}}_1 \ldots \forall\vec{\mathbf{x}}_n\exists\vec{\mathbf{y}}_n.\psi$ be an SF sentence of positive degree ∂_φ in standard form. Let $\mathcal{L}_\varphi(\vec{\mathbf{y}})$ denote the set of all literals in φ that contain at least one variable $y \in \vec{\mathbf{y}} := \vec{\mathbf{y}}_1 \cup \ldots \cup \vec{\mathbf{y}}_n$. There exists a sentence $\varphi_{\mathrm{BSR}} = \exists\vec{\mathbf{z}}\,\exists\vec{\mathbf{u}}\,\forall\vec{\mathbf{v}}.\psi_{\mathrm{BSR}}$ in standard form with quantifier-free ψ_{BSR} that is equivalent to φ and contains at most $|\vec{\mathbf{z}}| + |\vec{\mathbf{y}}| \cdot \partial_\varphi \cdot \big(2^{\uparrow\partial_\varphi}(|\mathcal{L}_\varphi(\vec{\mathbf{y}})|)\big)^{\partial_\varphi}$ leading existential quantifiers.

Proof sketch. Let $\vec{\mathbf{x}} := \vec{\mathbf{x}}_1 \cup \ldots \cup \vec{\mathbf{x}}_n$. We transform φ into CNF and then apply Lemma 11 and the rules of miniscoping given in Lemma 2 to push all quantifier blocks inwards. Since the sets $\vec{\mathbf{x}}$ and $\vec{\mathbf{y}}$ are separated in φ, these operations can be performed in such a way that in the resulting formula φ' no universal quantifier lies within the scope of any existential quantifier (other than the ones in $\exists\vec{\mathbf{z}}$) and vice versa. After removing redundant parts from φ', the depth of nestings of existential quantifier blocks (interspersed with conjunctive connectives in φ''s syntax tree) can be upper bounded by ∂_φ. As a consequence, φ' contains at most $2^{\uparrow\partial_\varphi}(|\mathcal{L}_\varphi(\vec{\mathbf{y}})|)$ distinct subformulas that are of the form $\exists y.\psi'$ and do not lie within the scope of any quantifier. After further transformations, we obtain a formula $\varphi'' := \bigvee_k \big(\chi_k(\vec{\mathbf{x}}) \wedge \bigwedge_{r_k} \eta_{r_k}(\vec{\mathbf{y}})\big)$ where the r_k range over at most $2^{\uparrow\partial_\varphi}(|\mathcal{L}_\varphi(\vec{\mathbf{y}})|)$ indices. Moreover, every constituent $\bigwedge_{r_k} \eta_{r_k}$ in φ'' contains at most $|\vec{\mathbf{y}}| \cdot \sum_{k'=1}^{\partial_\varphi} \prod_{d=k'}^{\partial_\varphi} 2^{\uparrow d}(|\mathcal{L}_\varphi(\vec{\mathbf{y}})|)$ occurrences of existential quantifiers. Since these existential quantifiers distribute over the topmost disjunction when we move them outwards to the front of the sentence φ'', and since the universal quantifiers in the χ_k may also be moved back outwards, one can show that φ is equivalent to some BSR sentence with at most $|\vec{\mathbf{y}}| \cdot \partial_\varphi \cdot \big(2^{\uparrow\partial_\varphi}(|\mathcal{L}_\varphi(\vec{\mathbf{y}})|)\big)^{\partial_\varphi}$ leading existential quantifiers. □

Proposition 6 now entails that any satisfiable SF-sentence φ has a model of size at most

$$\mathrm{len}(\varphi) + \mathrm{len}(\varphi) \cdot \partial_\varphi \cdot \big(2^{\uparrow\partial_\varphi}(\mathrm{len}(\varphi))\big)^{\partial_\varphi} . \tag{1}$$

Theorem 13. Let k be any positive integer. The satisfiability problem for the class of SF sentences φ in standard form with degree $\partial_\varphi \leq k$ can be decided in nondeterministic k-fold exponential time.

Together with Proposition 10(i), this establishes the upper bounds depicted in Figure 1.

In cases where $\partial_\varphi = 1$, Expression (1) simplifies to $\mathrm{len}(\varphi) + \mathrm{len}(\varphi) \cdot 2^{\mathrm{len}(\varphi)}$. The syntactic class of sentences satisfying this property is called *strongly separated* in [17].

Definition 14 ([17])**.** Let $\varphi := \forall\vec{\mathbf{x}}_1\exists\vec{\mathbf{y}}_1 \ldots \forall\vec{\mathbf{x}}_n\exists\vec{\mathbf{y}}_n.\psi$ be an SF sentence and assume that ψ is quantifier free. We say that φ belongs to the *strongly separated fragment (SSF)* if and only if the sets $\vec{\mathbf{x}} := \vec{\mathbf{x}}_1 \cup \ldots \cup \vec{\mathbf{x}}_n$ and $\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_n$ are all pairwise separated in φ.

Since MFO and BSR sentences fall into this syntactic category, and since their decision problem is known to be NExpTime-hard, we obtain the following corollary.

Corollary 15. The satisfiability problem for SSF is NExpTime-complete.

Notice that the presented method can explain the asymptotic complexity of MFO-satisfiability and yields a reasonable upper bound on the size of small models of satisfiable MFO sentences. This works in spite of the fact that monadic sentences may contain arbitrarily nested alternating quantifiers. This is a considerable improvement compared to the methods used in [17].

Let φ be any SF sentence with the maximally possible degree $\partial_\varphi = n$, where n is the number of occurring $\forall\exists$-alternations. Then the upper bound shown in Expression (1) regarding the number of elements in small models fits the corresponding result entailed by Proposition 9. As one consequence, Theorem 13 in the present paper subsumes Theorem 17 in [17]. Moreover, Corollary 15 improves the double exponential upper bound on SSF-satisfiability given in Theorem 15 in [17]. Finally, it is worth noticing that all SF sentences with the quantifier prefix $\exists^*\forall^*\exists^*\forall^*$ belong to the strongly separated fragment. Hence, Corollary 15 subsumes Theorem 14 in [17]. The latter stipulates NExpTime-completeness of SF sentences with quantifier prefix $\exists^*\forall^*\exists^*$. Clearly, the refined analysis based on the degree of interaction of existential variables, rather than the number of quantifier alternations, yields significantly tighter results in many cases.

4.2 Lower bounds on the length of equivalent BSR formulas

Before we derive lower bounds on the time that is required to decide SF-satisfiability in the worst case, we establish lower bounds on the length of the results of the translation from SF into the BSR fragment.

Theorem 16. There is a class of SF sentences that are Horn and Krom such that for every positive integer n the class contains a sentence φ of degree $\partial_\varphi = n$ and with a length linear in n for which any equivalent BSR sentence contains at least $\sum_{k=1}^{n} 2^{\uparrow k}(n)$ leading existential quantifiers.

Proof sketch. Recall that $[n]$ abbreviates the set $\{1, \ldots, n\}$ and that $\mathfrak{P}S$ denotes the power set of a given set S. Let $n \geq 1$ be some positive integer. Consider the following first-order sentence in which the sets $\{x_1, \ldots, x_n\}$ and $\{y_1, \ldots, y_n\}$ are separated:

$$\varphi := \forall x_n \exists y_n \ldots \forall x_1 \exists y_1 . \bigwedge_{i=1}^{4n} \big(P_i(x_1, \ldots, x_n) \leftrightarrow Q_i(y_1, \ldots, y_n)\big) \ .$$

Notice that we change the orientation of the indices in the quantifier prefix in this proof.

In order to construct a particular model of φ, we inductively define the following sets: $\mathcal{S}_1 := \big\{S \subseteq [4n] \;\big|\; |S| = 2n\big\}$, $\mathcal{S}_{k+1} := \big\{S \in \mathfrak{P}\mathcal{S}_k \;\big|\; |S| = \frac{1}{2} \cdot |\mathcal{S}_k|\big\}$ for every k, $1 < k \leq n$. Hence, we observe

$|\mathcal{S}_1| = \binom{4n}{2n} \geq \big(\frac{4n}{2n}\big)^{2n} = 2^{2n}$,

$|\mathcal{S}_2| = \binom{|\mathcal{S}_1|}{|\mathcal{S}_1|/2} \geq \big(\frac{|\mathcal{S}_1|}{|\mathcal{S}_1|/2}\big)^{|\mathcal{S}_1|/2} \geq 2^{2^{2n-1}}$,

$\vdots$

$|\mathcal{S}_n| = \binom{|\mathcal{S}_{n-1}|}{|\mathcal{S}_{n-1}|/2} \geq 2^{2^{\vdots^{2^{2n-1}-1}}-1} \geq 2^{\uparrow n}(n+1)$.

We now define the structure $\mathcal{A}$ as follows:

- $\mathfrak{U}_\mathcal{A} := \bigcup_{k=1}^{n} \big\{\mathfrak{a}_S^{(k)}, \mathfrak{b}_S^{(k)} \;\big|\; S \in \mathcal{S}_k\big\}$,
- $P_i^\mathcal{A} := \big\{\langle \mathfrak{a}_{S_1}^{(1)}, \ldots, \mathfrak{a}_{S_n}^{(n)}\rangle \in \mathfrak{U}_\mathcal{A}^n \;\big|\; i \in S_1 \in S_2 \in \ldots \in S_n\big\}$ for $i = 1, \ldots, 4n$, and
- $Q_i^\mathcal{A} := \big\{\langle \mathfrak{b}_{S_1}^{(1)}, \ldots, \mathfrak{b}_{S_n}^{(n)}\rangle \in \mathfrak{U}_\mathcal{A}^n \;\big|\; i \in S_1 \in S_2 \in \ldots \in S_n\big\}$ for $i = 1, \ldots, 4n$.

One can easily show that $\mathcal{A}$ is a model of φ. Moreover, employing a game-theoretic argument, one can show the following property:

($*$) the substructure induced by $\mathcal{A}$'s domain after removing at least one of the $\mathfrak{b}_S^{(k)}$ does not satisfy φ.

We know that $\mathfrak{U}_{\mathcal{A}}$ contains at least $\sum_{k=1}^{n} 2^{\uparrow k}(n)$ elements of the form $\mathfrak{b}_S^{(k)}$.

Using $(*)$ and the substructure lemma, one can argue that any BSR sentence φ_* that is semantically equivalent to φ must contain at least $\sum_{k=1}^{n} 2^{\uparrow k}(n)$ leading existential quantifiers.

The key idea is that φ_*, which is satisfied by $\mathcal{A}$, must contain one existential quantifier for each and every $\mathfrak{b}_S^{(k)}$. Otherwise, there would be one $\mathfrak{b}_S^{(k)}$, call it $\mathfrak{b}_*$, such that we could remove $\mathfrak{b}_*$ from $\mathcal{A}$'s domain and any tuple $\langle \ldots, \mathfrak{b}_*, \ldots \rangle$ from the sets $Q_i^{\mathcal{A}}$, and the resulting structure would then still be a model of φ_*. But this would contradict $(*)$.

Since every atom $Q_i(y_1, \ldots, y_n)$ contains n variables from existential quantifier blocks that are separated by universal ones, the degree ∂_φ of φ is n. Moreover, φ can easily be transformed into a CNF that is Horn and Krom at the same time.

Hence, the theorem holds. □

Theorem 16 entails that there is no elementary upper bound on the length of the BSR sentences that result from an equivalence-preserving transformation of SF sentences into BSR. On the other hand, by Lemma 12, there is an elementary upper bound, if we only consider SF sentences up to a certain degree.

5 Lower bounds on the computational complexity of SF-satisfiability

In this section we establish lower bounds on the worst-case time complexity of SF-satisfiability. Our arguments will be based on a particular form of bounded domino (or tiling) problems developed by Grädel (see [6] and [3], Section 6.1.1). By $\mathbb{Z}_t$ we denote the set of integers $\{0, \ldots, t-1\}$ for any positive $t \geq 1$.

Definition 17 ([3], Definition 6.1.1)**.** A *domino system* $\mathfrak{D} := \langle \mathcal{D}, \mathcal{H}, \mathcal{V} \rangle$ is a triple where $\mathcal{D}$ is a finite set of tiles and $\mathcal{H}, \mathcal{V} \subseteq \mathcal{D} \times \mathcal{D}$ are binary relations determining the allowed horizontal and vertical neighbors of tiles, respectively. Consider the torus $\mathbb{Z}_t^2 := \mathbb{Z}_t \times \mathbb{Z}_t$ and let $\overline{D} := D_0 \ldots D_{n-1}$ be a word over $\mathcal{D}$ of length $n \leq t$. The letters of $\overline{D}$ represent tiles. We say that $\mathfrak{D}$ *tiles the torus* $\mathbb{Z}_t^2$ *with initial condition* $\overline{D}$ if and only if there exists a mapping $\tau : \mathbb{Z}_t^2 \to \mathcal{D}$ such that for every $\langle x, y \rangle \in \mathbb{Z}_t^2$ the following conditions hold, where "+1" denotes increment modulo t.

(a) If $\tau(x, y) = D$ and $\tau(x+1, y) = D'$, then $\langle D, D' \rangle \in \mathcal{H}$.

(b) If $\tau(x, y) = D$ and $\tau(x, y+1) = D'$, then $\langle D, D' \rangle \in \mathcal{V}$.

(c) $\tau(i, 0) = D_i$ for $i = 0, \ldots, n-1$.

Definition 18 ([3], Definition 6.1.5)**.** Let $T : \mathbb{N} \to \mathbb{N}$ be a function and let $\mathfrak{D} := \langle \mathcal{D}, \mathcal{H}, \mathcal{V} \rangle$ be a domino system. The problem DOMINO$(\mathfrak{D}, T(n))$ is the set of those words $\overline{D}$ over the alphabet $\mathcal{D}$ for which $\mathfrak{D}$ tiles $\mathbb{Z}^2_{T(|\overline{D}|)}$ with initial condition $\overline{D}$.

Domino problems provide a convenient way of deriving lower bounds via reductions. Suppose we are given some well-behaved time bound $T(n)$ that grows sufficiently fast. Further assume there is a reasonable translation from DOMINO$(\mathfrak{D}, T(n))$ into some problem $\mathcal{L}$ where the length of the results is upper bounded by a function $g(n)$. It follows that the time required to solve the hardest instances of $\mathcal{L}$ lies in $\Omega\big(T(h(n))\big)$, where $h(n)$ may be conceived as an inverse of $g(n)$ from an asymptotic point of view. The next proposition formalizes this observation.

Proposition 19 ([3], Theorem 6.1.8)**.** Let $T : \mathbb{N} \to \mathbb{N}$ be a time-constructible function with $T(c'n)^2 \in o(T(n))$ for some constant $c' > 0$ and let $\mathcal{L}$ be a problem such that for every domino system $\mathfrak{D}$ we have DOMINO$(\mathfrak{D}, T(n)) \leq_{g(n)} \mathcal{L}$, i.e. DOMINO$(\mathfrak{D}, T(n))$ is *polynomially reducible* to $\mathcal{L}$ via length order g(n) (cf. Definition 6.1.7 in [3]). Moreover, let $h : \mathbb{N} \to \mathbb{N}$ be a function such that $h(d \cdot g(n)) \in \mathcal{O}(n)$ for any positive d. There exists a positive constant $c > 0$ such that $\mathcal{L} \notin \text{NTIME}(T(c \cdot h(n)))$.

Subsections 5.1 and 5.2 are devoted to the purpose of outlining the following reductions.

Lemma 20.

(i) Fix some positive integer $k > 0$ and let $\mathfrak{D}$ be an arbitrary domino system. Let $\mathrm{Sat}(\mathrm{SF}_{\partial\leq k})$ be the set containing all satisfiable SF sentences whose degree ∂ is at most k. We have $\mathrm{DOMINO}\big(\mathfrak{D}, 2^{\uparrow k}(n)\big) \leq_{n\cdot\log n} \mathrm{Sat}(\mathrm{SF}_{\partial\leq k})$.

(ii) Fix some positive integer $m > 1$ and let $\mathfrak{D}$ be an arbitrary domino system. Let $\mathrm{Sat}(\mathrm{SF})$ be the set containing all satisfiable SF sentences. We have $\mathrm{DOMINO}(\mathfrak{D}, 2^{\uparrow n}(m)) \leq_{n^2\cdot\log n} \mathrm{Sat}(\mathrm{SF})$.

Having these reduction results at hand, Proposition 19 implies the sought lower bounds on SF-satisfiability for classes of sentences with bounded degree and the class of unbounded SF sentences.

Theorem 21. There are positive constants $c, d > 0$ for which

$$\mathrm{Sat}(\mathrm{SF}_{\partial\leq k}) \notin \mathrm{NTIME}\big(2^{\uparrow k}(cn/\log n)\big)$$

and

$$\mathrm{Sat}(\mathrm{SF}) \notin \mathrm{NTIME}\big(2^{\uparrow d\cdot\sqrt{n/\log n}}(2)\big)\ .$$

These lower bounds also hold if we do not allow equality in SF, see Section 5.3. The remainder of Section 5 is devoted to the formalization of sufficiently large tori in SF and to the translation from a given domino system $\mathfrak{D} = \langle \mathcal{D}, \mathcal{H}, \mathcal{V}\rangle$ (for nonempty $\mathcal{D}, \mathcal{H}, \mathcal{V}$) plus an initial condition $\overline{D}$ into an SF sentence φ such that φ is satisfiable if and only if $\overline{D} \in \mathrm{DOMINO}(\mathfrak{D}, T_i(|\overline{D}|))$ with $T_1(n) = 2^{\uparrow\kappa}(n)$ for any given $\kappa > 0$ and $T_2(n) = 2^{\uparrow n}(\mu)$ for any given $\mu > 1$.

5.1 Enforcing a large domain

The following description gives a somewhat simplified view. Technical details will follow.

A crucial part in the reduction is that a grid of size $t \times t$ has to be encoded, where t defines the required computing time and we assume $t := 2^{\uparrow\kappa}(\mu)$ for positive integers κ and $\mu > 1$ that we consider as parameters of the construction.

Every point p on the grid is represented by a pair $p = \langle x, y\rangle$, where each of the coordinates x and y is represented by a bit string of length $\log\big(2^{\uparrow\kappa}(\mu)\big) = 2^{\uparrow\kappa-1}(\mu)$. Given a bit string $\overline{b}$, we represent the j-th bit b_j by the truth value of the atom $J(\underline{\kappa}, \overline{b}, j)$, where $\underline{\kappa}$ is the constant used to address the topmost level of a hierarchy of $\kappa + 1$ sets of indices. The crux of our approach is that we have to enforce the existence of sufficiently many indices j, namely $2^{\uparrow\kappa-1}(\mu)$ many, to address the single bits of $\overline{b}$. Again, we address each of these indices as a bit string, this time of length $2^{\uparrow\kappa-2}(\mu)$.

Thus, we proceed in an inductive fashion, building up a hierarchy of indices with $\kappa + 1$ levels. The lowest level, level zero, is inhabited by μ indices, which we represent as constants with distinct values. For every $\ell \geq 1$ any index j on the ℓ-th level is represented by a bit string consisting of $2^{\uparrow\ell-1}(\mu)$ bits, i.e. the ℓ-th level of the index hierarchy contains $2^{\uparrow\ell}(\mu)$ indices. The i-th bit of an ℓ-th-level index j corresponds to the truth value of the atom $J(\underline{\ell}, j, i)$.

Example 22. Assume $\mu = 2$ and $\kappa = 3$.

index level	set of indices	number of indices
0	$\{\mathfrak{c}_1, \mathfrak{c}_2\}$	2
1	$\{00, 01, 10, 11\}$	4
2	$\{0000, 0001, \ldots, 1111\}$	16
3	$\{0,1\}^{16}$	65536

On every index level, the bits of one index are indexed by the indices from the previous level. We illustrate this for the word 1010 on all levels from 2 down to 0. The bits of 1010 on level two are indexed by bit strings from level one, each of them having a length of two. The bits of the indices of level one are themselves indexed by objects of level zero which are some values $\mathfrak{c}_1, \mathfrak{c}_2$ assigned

to the constants c_1, c_2. To improve readability, we connect the bits of words by dashes.

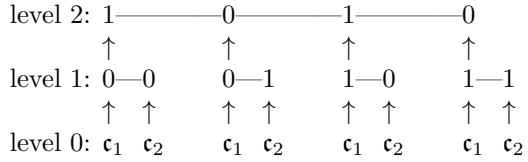
level 2: 1 0 1 0
level 1: 0 0 0 1 1 0 1 1
level 0: $\mathfrak{c}_1$ $\mathfrak{c}_2$ $\mathfrak{c}_1$ $\mathfrak{c}_2$ $\mathfrak{c}_1$ $\mathfrak{c}_2$ $\mathfrak{c}_1$ $\mathfrak{c}_2$

For technical reasons the number of indices per level grows slightly slower than described above (cf. Lemma 24). The described index hierarchies can be encoded by SF formulas with the quantifier prefix $\exists^*(\forall\exists)^\kappa$ that have a length that is polynomial in κ and μ. We use the following constant and predicate symbols with the indicated meaning:

$\underline{0}, \underline{1}, \ldots, \underline{\kappa}$	constants denoting the levels from 0 to κ,
$c_1, \ldots, c_\mu$	denote the indices at level 0,
$d_1, \ldots, d_\kappa$	d_ℓ is the min. index at level ℓ,
$e_1, \ldots, e_\kappa$	e_ℓ is the max. index at level ℓ,
$L(\underline{\ell}, j)$	index j belongs to level ℓ,
$\mathrm{MinIdx}(\underline{\ell}, j)$	j is a min. index at level ℓ,
$\mathrm{MaxIdx}(\underline{\ell}, j)$	j is a max. index at level ℓ,
$J(\underline{\ell}, j, i, b)$	the i-th bit of the index j at level ℓ is b,
$J^*(\underline{\ell}, j, i, b)$	$b = 1$ indicates that all the bits of the index j that are less significant than j's i-th bit are 1,
$\mathrm{Succ}(\underline{\ell}, j, j')$	j' is the successor index of j at level ℓ.

On every level we establish an ordering over the indices of that level. We use the usual ordering on natural numbers encoded in binary. Moreover, we formalize the usual successor relation on these numbers by the predicate Succ.

One difficulty that we encounter is that we cannot assert the existence of successors simply by adding $\forall j \exists j'. \mathrm{Succ}(\underline{\ell}, j, j')$, as j and j' would not be separated. Therefore, we fall back on a trick: we start from the equivalent formula $\forall j \exists \widetilde{j} j'. j \approx \widetilde{j} \wedge \mathrm{Succ}(\underline{\ell}, \widetilde{j}, j')$, and replace the atom $j \approx \widetilde{j}$ by a subformula $\mathrm{eq}^\ell_{j,\widetilde{j}}$ in which j and $\widetilde{j}$ are separated and which expresses a certain similarity between j and $\widetilde{j}$ instead of identity. However, we specify the hierarchy of indices in a sufficiently strong way such that the similarity expressed by $\mathrm{eq}^\ell_{j,\widetilde{j}}$ actually coincides with identity.

Next, we formalize the described index hierarchies in $\mathrm{SF}_{\partial \le \kappa}$. Every formula is accompanied by a brief description of its purpose. We shall try to use as few non-Horn sentences as possible.

$$\psi_1 := \bigwedge_{\ell=0}^{\kappa} \bigwedge_{\substack{\ell' = 0 \\ \ell' \neq \ell}}^{\kappa} \forall j.\ L(\underline{\ell}, j) \to \neg L(\underline{\ell'}, j)$$

Every index belongs to at most one level.

$$\psi_2 := \bigwedge_{\ell=0}^{\kappa} \big(\forall j.\ \mathrm{MinIdx}(\underline{\ell}, j) \to L(\underline{\ell}, j)\big) \quad \wedge \quad \big(\forall j j'.\ \mathrm{MinIdx}(\underline{\ell}, j) \to \neg \mathrm{Succ}(\underline{\ell}, j', j)\big)$$

A min. index of level ℓ belongs to level ℓ. A min. index does not have a predecessor.

$$\psi_3 := \bigwedge_{\ell=0}^{\kappa} \mathrm{MinIdx}(\underline{\ell}, d_\ell) \quad \wedge \quad \big(\forall j.\ \mathrm{MinIdx}(\underline{\ell}, j) \to j \approx d_\ell\big)$$

There is a unique min. index on every level.

$$\psi_4 := \bigwedge_{\ell=0}^{\kappa} \big(\forall j.\ \mathrm{MaxIdx}(\underline{\ell}, j) \to L(\underline{\ell}, j)\big) \quad \wedge \quad \big(\forall j j'.\ \mathrm{MaxIdx}(\underline{\ell}, j) \to \neg \mathrm{Succ}(\underline{\ell}, j, j')\big)$$

A max. index of level ℓ belongs to level ℓ. A max. index does not have a successor.

$$\psi_5 := \bigwedge_{\ell=0}^{\kappa} \mathrm{MaxIdx}(\underline{\ell}, e_\ell) \quad \wedge \quad \big(\forall j.\ \mathrm{MaxIdx}(\underline{\ell}, j) \to j \approx e_\ell\big)$$

There is a unique max. index on every level.

$$\psi_6 := \bigwedge_{\ell=0}^{\kappa} \forall j j'.\ \mathrm{Succ}(\underline{\ell}, j, j') \quad \to \quad L(\underline{\ell}, j) \wedge L(\underline{\ell}, j')$$

If j' is the successor of j at level ℓ, then both j and j' belong to level ℓ.

$$\psi_7 := \bigwedge_{\ell=0}^{\kappa} \forall j j' j''.\ \neg\mathrm{Succ}(\underline{\ell}, j, j) \quad \wedge \quad \big(\mathrm{Succ}(\underline{\ell}, j, j') \wedge \mathrm{Succ}(\underline{\ell}, j, j'') \to j' \approx j''\big)$$
$$\wedge \quad \big(\mathrm{Succ}(\underline{\ell}, j', j) \wedge \mathrm{Succ}(\underline{\ell}, j'', j) \to j' \approx j''\big)$$

The successor relation is irreflexive. Every index j has at most one successor and at most one predecessor.

$$\psi_8 := \mathrm{MinIdx}(\underline{0}, c_1) \wedge \mathrm{MaxIdx}(\underline{0}, c_\mu) \quad \wedge \quad \bigwedge_{i=1}^{\mu-1} \mathrm{Succ}(\underline{0}, c_i, c_{i+1})$$

At level zero we have the sequence $c_1, \ldots, c_\mu$ of successors, where c_1 is min. and c_μ max.

$$\psi_9 := \bigwedge_{\ell=1}^{\kappa} \forall j j' i.\ \mathrm{Succ}(\underline{\ell}, j, j') \wedge L(\underline{\ell-1}, i) \quad \to \quad \Big(\big(J^*(\underline{\ell}, j, i, 1) \wedge J(\underline{\ell}, j, i, 1) \to J(\underline{\ell}, j', i, 0)\big)$$
$$\wedge \big(J^*(\underline{\ell}, j, i, 1) \wedge J(\underline{\ell}, j, i, 0) \to J(\underline{\ell}, j', i, 1)\big)$$
$$\wedge \big(J^*(\underline{\ell}, j, i, 0) \wedge J(\underline{\ell}, j, i, 1) \to J(\underline{\ell}, j', i, 1)\big)$$
$$\wedge \big(J^*(\underline{\ell}, j, i, 0) \wedge J(\underline{\ell}, j, i, 0) \to J(\underline{\ell}, j', i, 0)\big)\Big)$$

Define what it means to be a successor at level ℓ, $\ell > 0$, in terms of the binary increment operation modulo $2^{\uparrow\ell}(\mu)$.

$$\psi_{10} := \bigwedge_{\ell=1}^{\kappa} \forall j i.\ \mathrm{MinIdx}(\underline{\ell}, j) \wedge L(\underline{\ell-1}, i) \quad \to \quad J(\underline{\ell}, j, i, 0)$$

All bits of a minimal index j are 0.

$$\psi_{11} := \bigwedge_{\ell=1}^{\kappa} \forall j i.\ \mathrm{MaxIdx}(\underline{\ell}, j) \wedge \mathrm{MaxIdx}(\underline{\ell-1}, i) \quad \to \quad J(\underline{\ell}, j, i, 1)$$

Define what it means to be max. (part 1): the most significant bit is 1.

$$\psi_{12} := \bigwedge_{\ell=1}^{\kappa} \forall j i.\ L(\underline{\ell}, j) \wedge \mathrm{MaxIdx}(\underline{\ell-1}, i) \wedge J(\underline{\ell}, j, i, 1) \quad \to \quad \mathrm{MaxIdx}(\underline{\ell}, j)$$

Define what it means to be max. (part 2): any index with 1 as its most significant bit is max.

$$\psi_{13} := \bigwedge_{\ell=1}^{\kappa} \forall j i.\ L(\underline{\ell}, j) \wedge L(\underline{\ell-1}, i) \quad \to \quad \big(J(\underline{\ell}, j, i, 0) \to \neg J(\underline{\ell}, j, i, 1)\big)$$
$$\wedge \big(J^*(\underline{\ell}, j, i, 0) \to \neg J^*(\underline{\ell}, j, i, 1)\big)$$

No bit of an index is 0 and 1 at the same time. An analogous requirement is stipulated for J^*.

$$\psi_{14} := \bigwedge_{\ell=1}^{\kappa} \forall j i.\ L(\underline{\ell}, j) \wedge \mathrm{MinIdx}(\underline{\ell-1}, i) \quad \to \quad J^*(\underline{\ell}, j, i, 1)$$

$J^*(\underline{\ell}, j, d_{\ell-1}, 1)$ holds for every index j.

$$\psi_{15} := \bigwedge_{\ell=1}^{\kappa} \forall j i i'.\ L(\underline{\ell}, j) \wedge \mathrm{Succ}(\underline{\ell-1}, i, i') \ \rightarrow\ \big(J^*(\underline{\ell}, j, i', 1) \leftrightarrow \big(J^*(\underline{\ell}, j, i, 1) \wedge J(\underline{\ell}, j, i, 1)\big)\big) \wedge \big(J(\underline{\ell}, j, i, 0) \rightarrow J^*(\underline{\ell}, j, i', 0)\big) \wedge \big(J^*(\underline{\ell}, j, i, 0) \rightarrow J^*(\underline{\ell}, j, i', 0)\big)$$

Define the semantics of J^* as indicating that all bits strictly less significant than the i-th bit are 1.

$$\mathrm{eq}^1_{j,\widetilde{j}} := L(\underline{1}, j) \wedge L(\underline{1}, \widetilde{j}) \ \wedge\ \bigwedge_{i=1}^{\mu} \big(J(\underline{1}, j, c_i, 0) \leftrightarrow J(\underline{1}, \widetilde{j}, c_i, 0)\big) \wedge \big(J(\underline{1}, j, c_i, 1) \leftrightarrow J(\underline{1}, \widetilde{j}, c_i, 1)\big)$$

Base case of equality of indices.

$$\mathrm{eq}^\ell_{j,\widetilde{j}} := L(\underline{\ell}, j) \wedge L(\underline{\ell}, \widetilde{j}) \wedge \forall i.\ L(\underline{\ell-1}, i) \ \rightarrow\ \exists \widetilde{i}.\ L(\underline{\ell-1}, \widetilde{i}) \wedge \mathrm{eq}^{\ell-1}_{i,\widetilde{i}} \wedge \big(J(\underline{\ell}, j, i, 0) \leftrightarrow J(\underline{\ell}, \widetilde{j}, \widetilde{i}, 0)\big) \wedge \big(J(\underline{\ell}, j, i, 1) \leftrightarrow J(\underline{\ell}, \widetilde{j}, \widetilde{i}, 1)\big)$$

Inductive case of equality of indices for $\ell > 1$.

$$\psi_{16} := \bigwedge_{\ell=1}^{\kappa} \forall j i.\ L(\underline{\ell}, j) \wedge \mathrm{MaxIdx}(\underline{\ell-1}, i) \wedge J(\underline{\ell}, j, i, 0) \ \rightarrow\ \exists \widetilde{j}\, \widetilde{j}'.\ \mathrm{eq}^\ell_{j,\widetilde{j}} \wedge \mathrm{Succ}(\underline{\ell}, \widetilde{j}, \widetilde{j}')$$

For every index at level ℓ that is not maximal, i.e. whose most significant bit is 0, there exists a successor index.

Until now, we have only introduced sentences that can easily be transformed into SF sentences in Horn form, existential variables are separated from universal ones, as all quantifiers occur with positive polarity, and consequents of implications are (conjunctions of) literals.

Regarding the length of the above sentences, we observe the following:

- $\mathrm{len}(\psi_1) \in \mathcal{O}(\kappa^2 \log \kappa)$
- $\mathrm{len}(\psi_2), \ldots, \mathrm{len}(\psi_7), \mathrm{len}(\psi_9), \ldots, \mathrm{len}(\psi_{15}) \in \mathcal{O}(\kappa \log \kappa)$
- $\mathrm{len}(\psi_8) \in \mathcal{O}\big(\mu(\log \kappa + \log \mu)\big)$
- $\mathrm{len}\big(\mathrm{eq}^1_{j,j'}\big) \in \mathcal{O}\big(\mu(\log \kappa + \log \mu)\big)$
- $\mathrm{len}\big(\mathrm{eq}^\ell_{j,j'}\big) \in \mathcal{O}(\log \kappa) + \mathrm{len}\big(\mathrm{eq}^1_{j,j'}\big)$
- $\mathrm{len}(\psi_{16}) \in \mathcal{O}\big(\kappa^2 \log \kappa + \kappa\mu(\log \kappa + \log \mu)\big)$

In total, this yields $\mathrm{len}(\psi_1 \wedge \ldots \wedge \psi_{16}) \in \mathcal{O}\big(\kappa^2 \log \kappa + \kappa\mu(\log \kappa + \log \mu)\big)$.

The following three sentences do not produce Horn formulas when transformed into CNF. They serve the purpose of removing spurious elements from the model. In particular, χ_3 is essential to enforce large models for $\kappa \geq 2$.

$$\chi_1 := \forall j.\ L(\underline{0}, j) \rightarrow \bigvee_{i=1}^{\mu} j \approx c_i$$

On level 0 there are no indices but $c_1, \ldots, c_\mu$.

$$\chi_2 := \bigwedge_{\ell=1}^{\kappa} \forall j i.\ L(\underline{\ell}, j) \wedge L(\underline{\ell-1}, i) \ \rightarrow\ J(\underline{\ell}, j, i, 0) \vee J(\underline{\ell}, j, i, 1)$$

We stipulate totality for the predicate J.

$$\chi_3 := \bigwedge_{\ell=1}^{\kappa} \forall j j'.\ L(\underline{\ell}, j) \wedge L(\underline{\ell}, j') \rightarrow \exists \widetilde{j}\, \widetilde{j}'.\ \mathrm{eq}^\ell_{j,\widetilde{j}} \wedge \mathrm{eq}^\ell_{j',\widetilde{j}'} \wedge \Big(\big(\forall \widetilde{i}.\ L(\underline{\ell-1}, \widetilde{i}) \rightarrow \big(J(\underline{\ell}, \widetilde{j}, \widetilde{i}, 0) \leftrightarrow J(\underline{\ell}, \widetilde{j}', \widetilde{i}, 0)\big)\big) \rightarrow j \approx j'\Big)$$

Two indices at the same level that agree on all of their bits are required to be identical.

Notice that χ_3 is (almost) an SF sentence, since the $\forall\widetilde{i}$ turns into a $\exists\widetilde{i}$ as soon as we bring the sentence into prenex normal form. Regarding the length of χ_1, χ_2, χ_3, we observe $\mathrm{len}(\chi_1) \in \mathcal{O}(\log\kappa + \mu\log\mu)$, $\mathrm{len}(\chi_2) \in \mathcal{O}(\kappa\log\kappa)$, and $\mathrm{len}(\chi_3) \in \mathcal{O}\big(\kappa^2\log\kappa + \kappa\mu(\log\kappa + \log\mu)\big)$. Hence, we overall have $\mathrm{len}(\chi_1 \wedge \chi_2 \wedge \chi_3) \in \mathcal{O}\big(\kappa^2\log\kappa + \kappa\mu(\log\kappa + \log\mu)\big)$.

Consider any model $\mathcal{A}$ of $\psi_1 \wedge \ldots \wedge \psi_{16} \wedge \chi_1 \wedge \chi_2 \wedge \chi_3$.

Definition 23. We define the following sets and relations: $\mathcal{I}_\ell := \big\{\mathfrak{a} \in \mathfrak{U}_\mathcal{A} \mid \mathcal{A}, [j{\mapsto}\mathfrak{a}] \models L(\underline{\ell}, j)\big\}$ for every $\ell = 0, \ldots, \kappa$; $\prec_\ell \subseteq \mathcal{I}_\ell \times \mathcal{I}_\ell$ for every $\ell = 0, \ldots, \kappa$ such that $\mathfrak{a} \prec_\ell \mathfrak{a}'$ holds if and only if $\mathcal{A}, [j{\mapsto}\mathfrak{a}, j'{\mapsto}\mathfrak{a}'] \models \mathrm{Succ}(\underline{\ell}, j, j')$.

Lemma 24. For every $\ell = 1, \ldots, \kappa$ we have $|\mathcal{I}_\ell| = p$ where $p := 2^{|\mathcal{I}_{\ell-1}|-1} + 1 = 2^{\uparrow\ell}(\mu - 1) + 1$. Moreover, there is a unique chain $\mathfrak{a}_1 \prec_\ell \ldots \prec_\ell \mathfrak{a}_p$ comprising all elements in $\mathcal{I}_\ell$.

Leaving out the non-Horn parts χ_1, χ_2, χ_3 renders the lemma invalid for $\ell > 1$. On the other hand, for $\kappa = 1$ the sentence $\psi_1 \wedge \ldots \wedge \psi_{16}$—which can be transformed into an equivalent Horn sentence—has only models $\mathcal{A}$ for which $\mathcal{I}_1$ contains at least $2^{\mu-1} + 1$ elements. Notice that this could be used to derive ExpTime-hardness of satisfiability for the class of Horn SF sentences of degree 1. But such lower bounds are already known for the Horn subfragments of MFO and of the BSR fragment, which are proper subsets of SF's Horn subfragment.

5.2 Formalizing a tiling of the torus

In order to formalize a given domino problem $\mathfrak{D} = \langle \mathcal{D}, \mathcal{H}, \mathcal{V}\rangle$ and an initial condition $\overline{D}$, we introduce the following constant and predicate symbols:

$H(x, y, x', y')$	$\langle x', y'\rangle$ is the horiz. neighbor of $\langle x, y\rangle$,
	i.e. $x' = x + 1 \pmod{2^{\uparrow\kappa}(\mu - 1) + 1}$ and $y' = y$,
$V(x, y, x', y')$	$\langle x', y'\rangle$ is the vert. neighbor of $\langle x, y\rangle$,
$\underline{D}(x, y)$	$\langle x, y\rangle$ is tiled with $D \in \mathcal{D}$,
$f_1, \ldots, f_{\lvert\overline{D}\rvert}$	constants addressing points $\langle 0, 0\rangle, \ldots, \langle\lvert\overline{D}\rvert - 1, 0\rangle$.

With the ideas we have seen when formalizing the index hierarchy, it is now fairly simple to formalize the torus. For instance, the following sentence makes sure that every point that is not on the "edge" of the torus has a horizontal neighbor.

$$\begin{aligned}\eta_3 := \forall xyi.\ & L(\underline{\kappa}, x) \wedge L(\underline{\kappa}, y) \wedge \mathrm{MaxIdx}(\underline{\kappa - 1}, i) \wedge J(\underline{\kappa}, x, i, 0)\\ & \rightarrow\ \exists\widetilde{x}\,\widetilde{y}\,\widetilde{x}'.\ \mathrm{eq}^\kappa_{x,\widetilde{x}} \wedge \mathrm{eq}^\kappa_{y,\widetilde{y}} \wedge \Big(\bigwedge_{D\in\mathcal{D}} \underline{D}(x, y) \leftrightarrow \underline{D}(\widetilde{x}, \widetilde{y})\Big) \wedge H(\widetilde{x}, \widetilde{y}, \widetilde{x}', \widetilde{y})\end{aligned}$$

The next sentence, on the other hand, makes sure that the rules of the domino system $\mathfrak{D}$ are obeyed.

$$\eta_{15} := \forall xx'y.\ H(x, y, x', y)\ \rightarrow \bigvee_{\langle D,D'\rangle\in\mathcal{H}} \underline{D}(x, y) \wedge \underline{D}'(x', y)$$

Proceeding this way, the formalization η of a domino system in SF requires a length in $\mathcal{O}\big(\widehat{n}\log\widehat{n}\big)$, where $\widehat{n} := \max\{\kappa, \mu, |\overline{D}|, |\mathcal{D}|^2\}$.

Lemma 25. Assume that $\mathcal{D}$, $\mathcal{H}$, and $\mathcal{V}$ are nonempty and let $\mathcal{A}$ be a model of the sentence $\psi_1 \wedge \ldots \wedge \psi_{16} \wedge \chi_1 \wedge \chi_2 \wedge \chi_3 \wedge \eta$. $\mathcal{A}$ induces a tiling τ of $\mathbb{Z}^2_t$ with initial condition $\overline{D} := D_1, \ldots, D_n$, where $t := 2^{\uparrow\kappa}(\mu - 1) + 1$.

5.3 Replacing the equality predicate

Since SF can express reflexivity, symmetry, transitivity, and congruence properties, it is easy to formulate an SF sentence without equality that is equisatisfiable to $\psi_1 \wedge \ldots \wedge \psi_{16} \wedge \chi_1 \wedge \chi_2 \wedge \chi_3 \wedge \eta$ and uses atoms $E(s, t)$ instead of $s \approx t$. In addition to replacing equational atoms as indicated, we

add the usual axioms concerning the fresh predicate symbol E. Overall, the additional formulas have a length that lies in $\mathcal{O}\big(\kappa \log \kappa + |\mathcal{D}| \log |\mathcal{D}|\big)$.

Consequently, the hardness result that we have obtained for SF with equality can be directly transferred to SF without equality. Moreover, notice that all the above formulas can be transformed into Horn form. Hence, one could also replace $\approx$ by E in hardness proofs for the Horn subfragment of SF.

6 Conclusion

We stress in this paper that an analysis of the computational complexity of satisfiability problems can greatly benefit from an analysis of how variables occur together in atoms instead of exclusively considering the number of occurring quantifier alternations. What we have not yet taken into account is the Boolean structure of sentences. This may widen the scope of our methods considerably and may moreover help understand where the hardness of satisfiability problems stems from.

Consider a quantified Boolean formula $\varphi := \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n . \psi$ with quantifier-free ψ. All satisfiable formulas of this shape together form a hard problem residing on the n-th level of the polynomial hierarchy. But what if, for instance, ψ has the form $\big(\bigwedge_i K_i\big) \wedge \big(\bigvee_j L_j\big)$, where the K_i and the L_j are literals and none of the existential variables in $\bigwedge_i K_i$ occurs in $\bigvee_j L_j$? Since Boolean variables cannot jointly occur in atoms, φ can be transformed into the equivalent formula $\exists \vec{\mathbf{y}}_1 \ldots \vec{\mathbf{y}}_n \forall \vec{\mathbf{x}}_1 \ldots \vec{\mathbf{x}}_n . \psi$ by application of the rules of miniscoping (cf. Lemma 2). Apparently, φ belongs to a class of QBF sentences that resides on the first level of the polynomial hierarchy rather than on the n-th.

Perhaps it is time to reconsider some of the definitions that are based on the shape of quantifier prefixes alone.

Acknowledgments

The author would like to thank the anonymous reviewers for valuable hints and suggestions.

References

[1] H. Andréka, I. Németi, and J. van Benthem. Modal Languages and Bounded Fragments of Predicate Logic. *J. Philosophical Logic*, 27(3):217–274, 1998.

[2] V. Bárány, B. ten Cate, and L. Segoufin. Guarded negation. *J. ACM*, 62(3):22, 2015.

[3] E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer, 1997.

[4] S. A. Cook. The complexity of theorem-proving procedures. In *Theory of Computing (STOC'71)*, pages 151–158, 1971.

[5] H. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic*. Springer, 1994.

[6] E. Grädel. On solvable cases of Hilbert's 'Entscheidungsproblem'. Habilitationsschrift, Universität Basel, 1990.

[7] E. Grädel. On the restraining power of guards. *J. Symb. Log.*, 64(4):1719–1742, 1999.

[8] E. Grädel, P. G. Kolaitis, and M. Y. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.

[9] N. D. Jones, Y. E. Lien, and W. T. Laaser. New problems complete for nondeterministic log space. *Mathematical systems theory*, 10(1):1–17, 1976.

[10] S. Kasif. On the parallel complexity of some constraint satisfaction problems. In *Artificial Intelligence (AAAI'86)*, volume 1, pages 349–353, 1986.

[11] E. Kieronski and A. Kuusisto. Complexity and expressivity of uniform one-dimensional fragment with equality. In *Mathematical Foundations of Computer Science (MFCS'14)*, pages 365–376, 2014.

[12] M. Mortimer. On Languages with Two Variables. *Mathematical Logic Quarterly*, 21(1):135–140, 1975.

[13] D. A. Plaisted. Complete problems in the first-order predicate calculus. *J. Comput. Syst. Sci.*, 29(1):8–35, 1984.

[14] I. Pratt-Hartmann, W. Szwast, and L. Tendera. Quine's fluted fragment is non-elementary. In *Computer Science Logic (CSL'16)*, pages 39:1–39:21, 2016.

[15] W. V. O. Quine. On the Limits of Decision. In *14th International Congress of Philosophy*, volume III, pages 57–62, 1969. An extended version appeared in W.V.Quine *Theories and Things*, Harvard University Press, 1981.

[16] W. V. O. Quine. The Variable. In W. V. O. Quine, editor, *The Ways of Paradox and other essays. Revised and enlarged edition.* Harvard University Press, 1976.

[17] T. Sturm, M. Voigt, and C. Weidenbach. Deciding First-Order Satisfiability when Universal and Existential Variables are Separated. In *Logic in Computer Science (LICS'16)*, pages 86–95, 2016.

[18] M. Voigt. A Fine-Grained Hierarchy of Hard Problems in the Separated Fragment. In *Logic in Computer Science (LICS'17)*, 2017. To appear.

Appendix

A Details regarding Section 3.2

Lemma 26.

(i) Satisfiability for the class of SF sentences φ without universal quantifiers and without equality can be decided nondeterministically in $\mathrm{poly}\big(\mathrm{len}(\varphi)\big)$ time.

(ii) Satisfiability for the class of SF-Horn sentences φ without universal quantifiers and without equality can be decided deterministically in $\mathrm{poly}\big(\mathrm{len}(\varphi)\big)$ time.

(iii) Satisfiability for the class of SF-Krom sentences φ without universal quantifiers and without equality can be decided nondeterministically using space that is logarithmic in $\mathrm{len}(\varphi)$.

Proof. We show that we can reduce the above SF-satisfiability problems to the respective satisfiability problems for propositional logic and vice versa.

Let φ be a first-order sentence without non-constant function symbols that does not contain any universal quantifiers. Skolemization of all its existential quantifiers leads to the equisatisfiable ground sentence φ_{gnd} in which every atom has the shape $P(c_1, \ldots, c_m)$. Let $A_1, \ldots, A_k$ be a complete enumeration of all the atoms—without duplicates—that occur in φ_{gnd}. Let $q_1, \ldots, q_k$ be a list of pairwise distinct propositional variables. We construct the propositional formula ϕ from φ_{gnd} by replacing every atom A_i with q_i. Clearly, any model of $\mathcal{A}$ of φ_{gnd} induces a model $\mathcal{B}$ of ϕ: $\mathcal{B} \models q_i$ if and only if $\mathcal{A} \models A_i$. Conversely, any model $\mathcal{B}'$ of ϕ induces a Herbrand model $\mathcal{A}'$ of φ: $\mathcal{A}' \models A_i$ if and only if $\mathcal{B}' \models q_i$. Consequently, deciding satisfiability of φ can be reduced to deciding satisfiability of ϕ. Moreover, we observe the following properties.

(a) $\mathrm{len}(\phi) \leq \mathrm{len}(\varphi)$.

(b) If φ is a Horn formula, then ϕ is Horn.

(c) If φ is a Krom formula, then ϕ is Krom. □

Lemma 27. There is an effective translation T from the class of SF sentences with equality but without universal quantifiers into the class of ground first-order sentences without equality such that for every φ from the former class we have that

(a) every model of $T(\varphi)$ is a model of φ,

(b) every model of φ can be extended to a model of $T(\varphi)$ over the same domain,

(c) $\mathrm{len}(T(\varphi)) \in \mathcal{O}\big(\mathrm{len}(\varphi)^3\big)$,

(d) If φ is Horn, then $T(\varphi)$ is Horn as well.

Proof. We describe the translation T informally. Let φ be some SF sentence that contains equality but no universal quantifiers. Let φ_{Sk} be the result of Skolemizing the leading existential quantifier block in φ. Notice that φ_{Sk} neither contains quantifiers nor non-constant function symbols. Let E be a binary predicate symbol that does not occur in φ_{Sk}. We construct the following ground formulas:

$$\psi_{\mathrm{refl}} := \bigwedge_{c \in \mathrm{consts}(\varphi_{\mathrm{Sk}})} E(c,c) \ ,$$

$$\psi_{\mathrm{symm}} := \bigwedge_{c,d \in \mathrm{consts}(\varphi_{\mathrm{Sk}})} \big(E(c,d) \rightarrow E(d,c)\big) \ ,$$

$$\psi_{\mathrm{trans}} := \bigwedge_{c,d,e \in \mathrm{consts}(\varphi_{\mathrm{Sk}})} \big(E(c,d) \wedge E(d,e) \rightarrow E(c,e)\big) \ .$$

Let ψ_{cong} be the conjunction of all ground formulas of the form $E(c_1, d_1) \wedge \ldots \wedge E(c_m, d_m) \wedge P(c_1, \ldots, c_m) \rightarrow P(d_1, \ldots, d_m)$ where $c_1, d_1, \ldots, c_m, d_m \in \text{consts}(\varphi_{\text{Sk}})$ and P is an m-ary predicate symbol in φ_{Sk}. We write ψ'_{cong} to denote the restriction of ψ_{cong} to formulas $E(c_1, d_1) \wedge \ldots \wedge E(c_m, d_m) \wedge P(c_1, \ldots, c_m) \rightarrow P(d_1, \ldots, d_m)$ whose constituents $P(c_1, \ldots, c_m)$ and $P(d_1, \ldots, d_m)$ actually occur in φ_{Sk}.

Let φ' be the result of replacing every equation $c \approx d$ in φ_{Sk} with the atom $E(c, d)$.

Claim I: φ_{Sk} is satisfiable if and only if $\varphi' \wedge \varphi_{\text{refl}} \wedge \varphi_{\text{symm}} \wedge \varphi_{\text{trans}} \wedge \varphi_{\text{cong}}$ is satisfiable.

Proof: Let $\mathcal{A}$ be a model of φ_{Sk}. By the Substructure Lemma, we may assume that $\mathcal{A}$'s universe is $\{c^{\mathcal{A}} \mid c \in \text{consts}(\varphi_{\text{Sk}})\}$. We now construct a model $\mathcal{B} \models \varphi' \wedge \varphi_{\text{refl}} \wedge \varphi_{\text{symm}} \wedge \varphi_{\text{trans}} \wedge \varphi_{\text{cong}}$ from $\mathcal{A}$. We take over $\mathcal{A}$'s domain and its interpretation of the predicate symbols and constant symbols in φ_{Sk}. We define E's interpretation under $\mathcal{B}$ such that $E^{\mathcal{B}} := \{\langle \mathfrak{a}, \mathfrak{a} \rangle \mid \mathfrak{a} \in \mathfrak{U}_{\mathcal{A}}\}$. Hence, for all $c, d \in \text{consts}(\varphi_{\text{Sk}})$ we observe $\mathcal{B} \models E(c, d)$ if and only if $\mathcal{B} \models c \approx d$. Consequently, $\mathcal{B}$ must be a model of φ_{Sk} and $\varphi' \wedge \varphi_{\text{refl}} \wedge \varphi_{\text{symm}} \wedge \varphi_{\text{trans}} \wedge \varphi_{\text{cong}}$.

Let $\mathcal{B}'$ be a model of $\varphi' \wedge \varphi_{\text{refl}} \wedge \varphi_{\text{symm}} \wedge \varphi_{\text{trans}} \wedge \varphi_{\text{cong}}$. By the Substructure Lemma, we may assume that $\mathcal{B}'$'s universe is $\mathfrak{U}_{\mathcal{B}'} = \{c^{\mathcal{B}'} \mid c \in \text{consts}(\varphi')\}$. We now construct a model $\mathcal{A}' \models \varphi_{\text{Sk}}$ from $\mathcal{B}'$. Because of $\mathcal{B}' \models \varphi_{\text{refl}} \wedge \varphi_{\text{symm}} \wedge \varphi_{\text{trans}} \wedge \varphi_{\text{cong}}$, we know that $E^{\mathcal{B}'}$ is a congruence relation over $\mathfrak{U}_{\mathcal{B}'}$. We define the universe of $\mathcal{A}'$ to be the quotient set $\mathfrak{U}_{\mathcal{A}'} := \mathfrak{U}_{\mathcal{B}'} / E^{\mathcal{B}'}$. Moreover, we define $c^{\mathcal{A}'} := [c^{\mathcal{B}'}]_{E^{\mathcal{B}'}}$ for every $c \in \text{consts}(\varphi_{\text{Sk}})$. For every congruence class $[\mathfrak{a}]_{E^{\mathcal{B}'}}$ we know that two domain elements $d^{\mathcal{B}'}, e^{\mathcal{B}'} \in [\mathfrak{a}]_{E^{\mathcal{B}'}}$ are indistinguishable by the relations $P^{\mathcal{B}'}$ for which P occurs in φ'. Therefore, we can use the following definition for every m-ary predicate symbol P in φ' (including E): $P^{\mathcal{A}'} := \{\langle [\mathfrak{a}_1]_{E^{\mathcal{B}'}}, \ldots, [\mathfrak{a}_m]_{E^{\mathcal{B}'}} \rangle \mid \langle \mathfrak{a}_1, \ldots, \mathfrak{a}_m \rangle \in P^{\mathcal{B}'}\}$. This yields $\mathcal{A}' \models \varphi'$ and for all $c, d \in \text{consts}(\varphi')$ we observe $\mathcal{A}' \models E(c, d)$ if and only if $\mathcal{A}' \models c \approx d$. Hence, $\mathcal{A}' \models \varphi_{\text{Sk}}$. ◊

It now remains to show equisatisfiability of $\varphi' \wedge \varphi_{\text{refl}} \wedge \varphi_{\text{symm}} \wedge \varphi_{\text{trans}} \wedge \varphi_{\text{cong}}$ and $\varphi' \wedge \varphi_{\text{refl}} \wedge \varphi_{\text{symm}} \wedge \varphi_{\text{trans}} \wedge \varphi'_{\text{cong}}$. The direction from left to right is obvious.

Claim II: Any model $\mathcal{A} \models \varphi' \wedge \varphi_{\text{refl}} \wedge \varphi_{\text{symm}} \wedge \varphi_{\text{trans}} \wedge \varphi'_{\text{cong}}$ gives rise to a model $\mathcal{B} \models \varphi' \wedge \varphi_{\text{refl}} \wedge \varphi_{\text{symm}} \wedge \varphi_{\text{trans}} \wedge \varphi_{\text{cong}}$.

Proof: Let $P(c_1, \ldots, c_m)$ and $P(d_1, \ldots, d_m)$ be two atoms that occur in φ'. We observe the following property:

(∗) If $E^{\mathcal{A}}$ contains the pairs $\langle c_1, d_1 \rangle, \ldots, \langle c_m, d_m \rangle$, then $\mathcal{A} \models \psi'_{\text{cong}}$ entails that $\mathcal{A} \models P(c_1, \ldots, c_m)$ holds if and only if $\mathcal{A} \models P(d_1, \ldots, d_m)$ does.

We define $\mathcal{B}$ such that $\mathfrak{U}_{\mathcal{B}} := \mathfrak{U}_{\mathcal{A}}$, $E^{\mathcal{B}} := E^{\mathcal{A}}$, and $c^{\mathcal{B}} := c^{\mathcal{A}}$ for every $c \in \text{consts}(\varphi')$. Moreover, for every m-ary predicate symbol P occurring in φ' and every tuple $\langle \mathfrak{a}_1, \ldots, \mathfrak{a}_m \rangle \in \mathfrak{U}_{\mathcal{B}}$ we set $\langle \mathfrak{a}_1, \ldots, \mathfrak{a}_m \rangle \in P^{\mathcal{B}}$ if and only if there is some atom $P(c_1, \ldots, c_m)$ in φ' for which we have $\langle c_1^{\mathcal{B}}, \mathfrak{a}_1 \rangle, \ldots, \langle c_m^{\mathcal{B}}, \mathfrak{a}_m \rangle \in E^{\mathcal{B}}$. Due to (∗), we know that $\mathcal{B} \models \varphi'$ still holds. By construction of $\mathcal{B}$, we moreover observe $\mathcal{B} \models \varphi_{\text{refl}} \wedge \varphi_{\text{symm}} \wedge \varphi_{\text{trans}} \wedge \varphi_{\text{cong}}$. ◊

We set $T(\varphi) := \varphi' \wedge \varphi_{\text{refl}} \wedge \varphi_{\text{symm}} \wedge \varphi_{\text{trans}} \wedge \varphi'_{\text{cong}}$. It is easy to see that if φ is Horn then $T(\varphi)$ is Horn as well. Moreover, the length of $T(\varphi)$ is upper bounded by $k \cdot \big(\text{len}(\varphi') + |\text{consts}(\varphi_{\text{Sk}})|^3 + \text{len}(\varphi') \cdot |\text{At}(\varphi_{\text{Sk}})|^2\big)$ for some positive integer k, where $\text{At}(\varphi_{\text{Sk}})$ denotes the set of all non-equational atoms that occur in φ_{Sk}. □

Proposition 10.

(i) Satisfiability for the class of SF sentences without universal quantifiers is NP-complete.

(ii) Satisfiability for the class of SF-Horn sentences without universal quantifiers is P-complete.

(iii) Satisfiability for the class of SF-Krom sentences without universal quantifiers and without equality is NL-complete.

Proof. The membership in the respective complexity classes is settled in Lemmas 26 and 27. It remains to reduce the respective SAT problems to the corresponding satisfiability problems for SF.

Let ϕ' be some propositional sentence and let $q_1, \ldots, q_k$ be a complete list of all propositional variables occurring in ϕ' (without duplicates). Let $c_1, \ldots, c_k$ be pairwise distinct constant symbols. We construct the first-order sentence φ' from ϕ' by replacing every q_i by the atom $P(c_i)$. By similar arguments as we have used in the proof of Lemma 26, we can show that ϕ' is satisfiable if and only if φ' is satisfiable. Moreover, we observe the following properties.

(a) $\mathrm{len}(\varphi) \in \mathcal{O}\big(\mathrm{len}(\phi)\big)$.

(b) If ϕ is a Horn formula, then φ' is Horn.

(c) If ϕ is a Krom formula, then φ is Krom. □

B Details regarding Section 4

Lemma 11. Let I and K_i, $i \in I$, be sets that are finite, nonempty, and pairwise disjoint. The elements of these sets serve as indices. Let

$$\varphi := \exists \vec{\mathbf{y}}.\, \bigwedge_{i \in I} \Big(\chi_i(\vec{\mathbf{z}}) \vee \bigvee_{k \in K_i} \eta_k(\vec{\mathbf{y}}, \vec{\mathbf{z}}) \Big)$$

be some first-order formula where the χ_i and the η_k denote arbitrary subformulas that we treat as indivisible units in what follows. We say that $f : I \to \bigcup_{i \in I} K_i$ is a *selection function* if for every $i \in I$ we have $f(i) \in K_i$. We denote the set of all selection functions of this form by $\mathcal{F}$.

φ is equivalent to

$$\varphi' := \bigwedge_{\substack{S \subseteq I \\ S \neq \emptyset}} \Big(\bigvee_{i \in S} \chi_i(\vec{\mathbf{z}}) \Big) \vee \bigvee_{f \in \mathcal{F}} \Big(\exists \vec{\mathbf{y}}.\, \bigwedge_{i \in S} \eta_{f(i)}(\vec{\mathbf{y}}, \vec{\mathbf{z}}) \Big) \,.$$

Proof. For the sake of readability we sometimes reuse variables in different occurrences of quantifiers in this proof. Using distributivity of $\wedge$ over $\vee$, we transform φ into an equivalent disjunction of conjunctions:

$$\exists \vec{\mathbf{y}}. \bigvee_{\substack{\langle T,\, f \rangle \in \\ (\mathfrak{P} I) \times \mathcal{F}}} \Big(\bigwedge_{i \in T} \chi_i(\vec{\mathbf{z}}) \Big) \wedge \Big(\bigwedge_{i \in I \setminus T} \eta_{f(i)}(\vec{\mathbf{y}}, \vec{\mathbf{z}}) \Big) \,.$$

Since the existential quantifier block distributes over the topmost disjunction, we can move this block inwards and obtain the equivalent formula

$$\bigvee_{\substack{\langle T,\, f \rangle \in \\ (\mathfrak{P} I) \times \mathcal{F}}} \Big(\bigwedge_{i \in T} \chi_i(\vec{\mathbf{z}}) \Big) \wedge \Big(\exists \vec{\mathbf{y}}.\, \bigwedge_{i \in I \setminus T} \eta_{f(i)}(\vec{\mathbf{y}}, \vec{\mathbf{z}}) \Big) \,. \tag{2}$$

At this point, we employ distributivity of $\vee$ over $\wedge$ to transform this result into an equivalent conjunction of disjunctions $\varphi'' := \bigwedge_j \bigvee_\ell \psi_{j,\ell}$ in which for every index j and every pair $\langle T, f \rangle \in (\mathfrak{P} I) \times \mathcal{F}$ there is exactly one ℓ such that

- either $\psi_{j,\ell} = \chi_i$ for some $i \in T$
- or $\psi_{j,\ell} = \exists \vec{\mathbf{y}}.\, \bigwedge_{i \in I \setminus T} \eta_{f(i)}(\vec{\mathbf{y}}, \vec{\mathbf{z}})$.

In order to show that φ'' is semantically equivalent to φ', we prove the following claims.

Claim I: Every disjunction $\bigvee_\ell \psi_{j,\ell}$ in φ'' is subsumed by a disjunction

$$\psi'_S := \Big(\bigvee_{i\in S} \chi_i(\vec{\mathbf{z}})\Big) \vee \bigvee_{f\in\mathcal{F}} \Big(\exists\vec{\mathbf{y}}.\bigwedge_{i\in S} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big)$$

for some nonempty $S \subseteq I$.

Proof: Fix some index j and consider $\bigvee_\ell \psi_{j,\ell}$. We set $S := \big\{i \in I \;\big|\; \psi_{j,\ell} = \chi_i \text{ for some } \ell\big\}$. Consider the set $\overline{S} := I \setminus S$. By definition of S, we know that none of the χ_i with $i \in \overline{S}$ is a constituent of $\bigvee_\ell \psi_{j,\ell}$. For every selection function $f \in \mathcal{F}$ there is some disjunct

$$\Big(\bigwedge_{i\in\overline{S}} \chi_i(\vec{\mathbf{z}})\Big) \wedge \Big(\exists\vec{\mathbf{y}}.\bigwedge_{i\in I\setminus\overline{S}} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big)$$

in (2) of which we know that none of the χ_i in it has been picked as constituent of $\bigvee_\ell \psi_{j,\ell}$ when constructing φ''. Hence, due to the definition of φ'', there must be some ℓ_* such that $\psi_{j,\ell_*} = \exists\vec{\mathbf{y}}.\bigwedge_{i\in I\setminus\overline{S}} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})$, where $I \setminus \overline{S} = S$.

Consequently, $\bigvee_\ell \psi_{j,\ell}$ is subsumed by ψ'_S. ◊

Claim II: Each of the subsuming disjunctions ψ'_S in Claim I is indeed equivalent to some disjunction $\bigvee_\ell \psi_{j,\ell}$ in φ''.

Proof: Fix some nonempty $S_* \subseteq I$ and consider ψ'_{S_*}. We obtain the equivalent disjunction ψ_* from the disjuncts in (2) as follows. For every $T \subseteq I$ with nonempty $T \cap S_*$ we pick one of the χ_i with $i \in T \cap S_*$ as constituent of ψ_*. For every $T \subseteq I$ for which $T \cap S_*$ is empty and any $f \in \mathcal{F}$ we pick $\exists\vec{\mathbf{y}}.\bigwedge_{i\in I\setminus T} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})$ as constituent of ψ_*. Since S_* is nonempty, T must be a proper subset of I and thus $I \setminus T$ is also nonempty.

For every constituent of the form $\exists\vec{\mathbf{y}}.\bigwedge_{i\in T'} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})$ that belongs to the disjunction ψ_* we know that $S_* \subseteq T'$. Hence, ψ_* is of the form

$$\Big(\bigvee_{i\in S_*} \chi_i(\vec{\mathbf{z}})\Big) \vee \bigvee_j \bigvee_{f\in\mathcal{F}} \exists\vec{\mathbf{y}}.\Big(\bigwedge_{i\in S_*} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big) \wedge \Big(\bigwedge_{i\in S'_j} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big)$$

for certain sets $S'_j \subseteq I \setminus S_*$. Among the S'_j is, in particular, the empty set, originating from $T = I \setminus S_*$. In this case, we have $S'_j = \big(I \setminus T\big) \setminus S_* = \big(I \setminus (I \setminus S_*)\big) \setminus S_* = \emptyset$. Hence, we can equivalently transform ψ_* into

$$\begin{aligned}
&\Big(\bigvee_{i\in S_*} \chi_i(\vec{\mathbf{z}})\Big) \vee \bigvee_{f\in\mathcal{F}} \bigvee_j \exists\vec{\mathbf{y}}.\Big(\bigwedge_{i\in S_*} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big) \wedge \Big(\bigwedge_{i\in S'_j} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big) \\
&\models \Big(\bigvee_{i\in S_*} \chi_i(\vec{\mathbf{z}})\Big) \vee \bigvee_{f\in\mathcal{F}} \exists\vec{\mathbf{y}}.\bigvee_j \Big(\Big(\bigwedge_{i\in S_*} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big) \wedge \Big(\bigwedge_{i\in S'_j} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big)\Big) \\
&\models \Big(\bigvee_{i\in S_*} \chi_i(\vec{\mathbf{z}})\Big) \vee \bigvee_{f\in\mathcal{F}} \exists\vec{\mathbf{y}}.\Big(\bigwedge_{i\in S_*} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big) \vee \bigvee_j \Big(\bigwedge_{i\in S_*} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}}) \wedge \bigwedge_{\substack{i\,\in\, S'_j \\ S'_j \neq \emptyset}} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big) \\
&\models \Big(\bigvee_{i\in S_*} \chi_i(\vec{\mathbf{z}})\Big) \vee \bigvee_{f\in\mathcal{F}} \exists\vec{\mathbf{y}}.\Big(\bigwedge_{i\in S_*} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big) \vee \Big(\Big(\bigwedge_{i\in S_*} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big) \wedge \bigvee_j \bigwedge_{\substack{i\,\in\, S'_j \\ S'_j \neq \emptyset}} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big) .
\end{aligned}$$

By the absorption axiom of Boolean algebra, we finally obtain the equivalent disjunction

$$\Big(\bigvee_{i\in S_*} \chi_i(\vec{\mathbf{z}})\Big) \vee \bigvee_{f\in\mathcal{F}} \Big(\exists\vec{\mathbf{y}}.\bigwedge_{i\in S_*} \eta_{f(i)}(\vec{\mathbf{y}},\vec{\mathbf{z}})\Big) .$$

Thus, the claimed equivalence holds.

We have not yet explicitly explained why the first subformula $\bigvee_{i\in S_*}\chi_i(\vec{\mathbf{z}})$ of ψ_* covers S_* completely. But this is easy to see, when one takes the singleton sets $T=\{i\}$ for every $i\in S_*$ into account, for which we pick the χ_i as a constituents of ψ_*. ◊

This completes the proof of the lemma. □

Lemma 12. Let $\varphi := \exists\vec{\mathbf{z}}\forall\vec{\mathbf{x}}_1\exists\vec{\mathbf{y}}_1\ldots\forall\vec{\mathbf{x}}_n\exists\vec{\mathbf{y}}_n.\psi$ be an SF sentence of positive degree ∂_φ in standard form. Let $\mathcal{L}_\varphi(\vec{\mathbf{y}})$ denote the set of all literals in φ that contain at least one variable $y\in\vec{\mathbf{y}} := \vec{\mathbf{y}}_1\cup\ldots\cup\vec{\mathbf{y}}_n$. There exists a sentence $\varphi_{\mathrm{BSR}} = \exists\vec{\mathbf{z}}\,\exists\vec{\mathbf{u}}\,\forall\vec{\mathbf{v}}.\psi_{\mathrm{BSR}}$ in standard form with quantifier-free ψ_{BSR} that is equivalent to φ and contains at most $|\vec{\mathbf{z}}|+|\vec{\mathbf{y}}|\cdot\partial_\varphi\cdot\big(2^{\uparrow\partial_\varphi}(|\mathcal{L}_\varphi(\vec{\mathbf{y}})|)\big)^{\partial_\varphi}$ leading existential quantifiers.

Proof. For convenience, we pretend that $\vec{\mathbf{z}}$ is empty. The argument works for nonempty $\vec{\mathbf{z}}$ as well. Let $\vec{\mathbf{x}} := \vec{\mathbf{x}}_1\cup\ldots\cup\vec{\mathbf{x}}_n$. We transform φ into an equivalent CNF formula of the form

$$\forall\vec{\mathbf{x}}_1\exists\vec{\mathbf{y}}_1\ldots\forall\vec{\mathbf{x}}_n\exists\vec{\mathbf{y}}_n.\bigwedge_{i\in I}\Big(\chi_i(\vec{\mathbf{x}})\vee\bigvee_{k\in K_i}L_k(\vec{\mathbf{y}})\Big)$$

where I and the K_i are finite, pairwise disjoint sets of indices, the subformulas χ_i are disjunctions of literals, and the L_k are literals. By Lemma 11, we can construct an equivalent formula of the form

$$\varphi' := \forall\vec{\mathbf{x}}_1\exists\vec{\mathbf{y}}_1\ldots\forall\vec{\mathbf{x}}_n.\bigwedge_{S\in\mathfrak{P}I\setminus\emptyset}\Big(\bigvee_{i\in S}\chi_i(\vec{\mathbf{x}})\Big)\vee\bigvee_{f\in\mathcal{F}}\Big(\exists\vec{\mathbf{y}}_n.\bigwedge_{i\in S}\eta_{f(i)}(\vec{\mathbf{y}})\Big)$$

where $\mathcal{F}$ is the set of all selection functions over the index sets K_i, $i\in I$. For the sake of readability we sometimes reuse variables in different occurrences of quantifiers in this proof. Applying ordinary miniscoping, we may move inward the universal quantifier block $\forall\vec{\mathbf{x}}_n$ and thus obtain

$$\varphi'' := \forall\vec{\mathbf{x}}_1\exists\vec{\mathbf{y}}_1\ldots\exists\vec{\mathbf{y}}_{n-1}.\bigwedge_{S\in\mathfrak{P}I\setminus\emptyset}\Big(\forall\vec{\mathbf{x}}_n.\bigvee_{i\in S}\chi_i(\vec{\mathbf{x}})\Big)\vee\bigvee_{f\in\mathcal{F}}\Big(\exists\vec{\mathbf{y}}_n.\bigwedge_{i\in S}\eta_{f(i)}(\vec{\mathbf{y}})\Big)\ .$$

We can now iterate these two steps in an alternating fashion until all quantifier blocks have been moved inwards in the described way. The constituents of the result $\varphi^{(3)} := \bigwedge_q\Big(\chi_q^{(3)}\vee\bigvee_p\eta_{qp}^{(3)}\Big)$ of this process have the form

$$\chi_q^{(3)} = \forall\vec{\mathbf{x}}_1.\bigvee_{\ell_1}\forall\vec{\mathbf{x}}_2.\bigvee_{\ell_2}\Big(\ldots\Big(\bigvee_{\ell_{n-1}}\forall\vec{\mathbf{x}}_n.\bigvee_{i\in S_{\ell_1,\ldots,\ell_{n-1}}}\chi_i(\vec{\mathbf{x}})\Big)\ldots\Big)$$

where the $S_{\ell_1,\ldots,\ell_n-1}$ are certain subsets of I, and

$$\eta_{qp}^{(3)} = \exists\vec{\mathbf{y}}_1.\bigwedge_{\ell_1'}\exists\vec{\mathbf{y}}_2.\bigwedge_{\ell_2'}\Big(\ldots\Big(\bigwedge_{\ell_{n-1}'}\exists\vec{\mathbf{y}}_n.\bigwedge_{k\in K_{\ell_1',\ldots,\ell_{n-1}'}}L_k(\vec{\mathbf{y}})\Big)\ldots\Big)$$

where the $K_{\ell_1',\ldots,\ell_{n-1}'}$ are certain subsets of $\bigcup_{i\in I}K_i$.

By definition of ∂_φ, we may assume that there is some positive integer m and a partition of the set $\vec{\mathbf{y}}$ into m nonempty subsets $Y_1,\ldots,Y_m$ such that $Y_1,\ldots,Y_m$ are all pairwise separated in φ and for every j, $1\le j\le m$, we have $\partial_{Y_j}\le\partial_\varphi$. Since the sets $Y_1,\ldots,Y_m$ are pairwise separated in φ, we may partition the set $\mathcal{L}_\varphi(\vec{\mathbf{y}})$ into subsets $\mathcal{L}_\varphi(Y_1),\ldots,\mathcal{L}_\varphi(Y_m)$ such that each $\mathcal{L}_\varphi(Y_j)$ contains exactly the literals in φ that contain at least one variable from Y_j. This means, we can rewrite every $\eta_{qp}^{(3)}$ into the form

$$\eta_{qp}^{(4)} = \exists\vec{\mathbf{y}}_1.\bigwedge_{\ell_1'}\exists\vec{\mathbf{y}}_2.\bigwedge_{\ell_2'}\Big(\ldots\Big(\bigwedge_{\ell_{n-1}'}\exists\vec{\mathbf{y}}_n.\bigwedge_{j\in[m]}\bigwedge_{k\in K^j_{\ell_1',\ldots,\ell_{n-1}'}}L_k(Y_j)\Big)\ldots\Big)$$

where the sets $K^1_{\ell'_1,\ldots,\ell'_{n-1}},\ldots,K^m_{\ell'_1,\ldots,\ell'_{n-1}}$ constitute a partition of $K_{\ell'_1,\ldots,\ell'_{n-1}}$. (Some of these parts may be empty.) We then observe the following equivalences.

$$\begin{aligned}
&\exists\vec{\mathbf{y}}_1.\bigwedge_{\ell'_1}\exists\vec{\mathbf{y}}_2.\bigwedge_{\ell'_2}\Big(\ldots\Big(\bigwedge_{\ell'_{n-1}}\exists\vec{\mathbf{y}}_n.\bigwedge_{j\in[m]}\bigwedge_{k\in K^j_{\ell'_1\ldots\ell'_{n-1}}}L_k(Y_j)\Big)\ldots\Big)\\
\models\ &\exists\vec{\mathbf{y}}_1.\bigwedge_{\ell'_1}\exists\vec{\mathbf{y}}_2.\bigwedge_{\ell'_2}\Big(\ldots\Big(\bigwedge_{\ell'_{n-1}}\bigwedge_{j\in[m]}\exists(\vec{\mathbf{y}}_n\cap Y_j).\bigwedge_{k\in K^j_{\ell'_1\ldots\ell'_{n-1}}}L_k(Y_j)\Big)\ldots\Big)\\
\models\ &\exists\vec{\mathbf{y}}_1.\bigwedge_{\ell'_1}\exists\vec{\mathbf{y}}_2.\bigwedge_{\ell'_2}\Big(\ldots\Big(\bigwedge_{j\in[m]}\bigwedge_{\ell''_{n-1}}\exists(\vec{\mathbf{y}}_n\cap Y_j).\bigwedge_{k\in K^j_{\ell'_1\ldots\ell'_{n-2}\ell''_{n-1}}}L_k(Y_j)\Big)\ldots\Big)\\
&\ \ \vdots\\
\models\ &\bigwedge_{j\in[m]}\exists(\vec{\mathbf{y}}_1\cap Y_j).\bigwedge_{\ell''_1}\exists(\vec{\mathbf{y}}_2\cap Y_j).\bigwedge_{\ell''_2}\Big(\ldots\Big(\bigwedge_{\ell''_{n-1}}\exists(\vec{\mathbf{y}}_n\cap Y_j).\bigwedge_{k\in K^j_{\ell''_1\ldots\ell''_{n-1}}}L_k(Y_j)\Big)\ldots\Big)
\end{aligned}$$

For every $\eta^{(4)}_{qp}$, we call the result of the above transformation $\eta^{(5)}_{qp}$. In cases where the set $\vec{\mathbf{y}}_i\cap Y_j$ is empty, the existential quantifier block vanishes. For every $j\in[m]$ there are at most ∂_{Y_j} nonempty sets $\vec{\mathbf{y}}_i\cap Y_j$. Hence, every $\eta^{(5)}_{qp}$ contains at most ∂_φ nested existential quantifier blocks that are separated by in-between conjunctive connectives in the syntax tree.

We obtain $\varphi^{(5)}$ from $\varphi^{(3)}$ by replacing every constituent $\eta^{(3)}_{qp}$ with the corresponding $\eta^{(5)}_{qp}$ after applying the idempotence axioms of Boolean Algebra exhaustively to remove redundant conjuncts.

Let $\kappa:=\max\big\{|\mathcal{L}_\varphi(Y_j)|\;\big|\;j\in[m]\big\}$. Due to the idempotence axioms, the following upper bounds can be shown inductively for any positive integer d, starting from $d=1$: Modulo idempotence, there are at most $2^{\uparrow d}(\kappa)$ formulas of the form

$$\exists(\vec{\mathbf{y}}_{i_1}\cap Y_j).\bigwedge_{\ell_1}\exists(\vec{\mathbf{y}}_{i_2}\cap Y_j).\bigwedge_{\ell_2}\Big(\ldots\Big(\bigwedge_{\ell_{d-1}}\exists(\vec{\mathbf{y}}_{i_d}\cap Y_j).\bigwedge_{k\in K'^{ij}_{\ell_1,\ldots,\ell_{d-1}}}L_k(Y_j)\Big)\ldots\Big)\,.$$

For the sentence $\varphi^{(5)}=\bigwedge_q\Big(\chi^{(3)}_q\vee\bigvee_p\eta^{(5)}_{qp}\Big)$ this means that it contains at most $2^{\uparrow\partial_\varphi}(\kappa)$ distinct subformulas (not occurrences thereof!) that are of the form $\exists y.\psi'$ and do not lie within the scope of any quantifier. We treat every such subformula $\exists y.\psi'$ and every subformula $\chi^{(3)}_q$ as indivisible unit and, employing distributivity of $\wedge$ over $\vee$, transform $\varphi^{(5)}$ into a disjunction of conjunctions $\varphi^{(6)}:=\bigvee_s\Big(\bigwedge_{r_1}\chi^{(6)}_{r_1}\wedge\bigwedge_{r_2}\eta^{(6)}_{r_2}\Big)$ where the $\chi^{(6)}_{r_1}$ have the same shape as the $\chi^{(3)}_q$, and the $\eta^{(6)}_{r_2}$ are of the form

$$\exists(\vec{\mathbf{y}}_{i_1}\cap Y_j).\bigwedge_{\ell_1}\exists(\vec{\mathbf{y}}_{i_2}\cap Y_j).\bigwedge_{\ell_2}\Big(\ldots\Big(\bigwedge_{\ell_{\partial_{Y_j}-1}}\exists(\vec{\mathbf{y}}_{i_{\partial_{Y_j}}}\cap Y_j).\bigwedge_{k\in K'^{j}_{\ell_1,\ldots,\ell_{\partial_{Y_j}-1}}}L_k(Y_j)\Big)\ldots\Big)$$

for a certain j and certain indices $i_1,\ldots,i_{\partial_{Y_j}}$ with $1\le i_1<\ldots<i_{\partial_{Y_j}}\le n$, all depending on r_2.

Due to previous observations, we know that, modulo idempotence, r_2 ranges over at most $2^{\uparrow\partial_\varphi}(\kappa)$ indices. Moreover, any ℓ_k in any $\eta^{(6)}_{r_2}$ ranges over at most $2^{\uparrow\partial_\varphi-k}(\kappa)$ indices. Consequently, every constituent $\bigwedge_{r_2}\eta^{(6)}_{r_2}$ in $\varphi^{(6)}$ contains at most $\max_{i,j}|\vec{\mathbf{y}}_i\cap Y_j|\cdot\sum_{k'=1}^{\partial_\varphi}\prod_{d=k'}^{\partial_\varphi}2^{\uparrow d}(\kappa)$ occurrences of existential quantifiers.

Since these existential quantifiers distribute over the topmost disjunction when we move them outwards to the front of the sentence $\varphi^{(6)}$, and since the universal quantifiers in the $\chi^{(6)}_q$ may also be moved back outwards, we have shown that φ is equivalent to some BSR sentence with at most $|\vec{\mathbf{y}}|\cdot\partial_\varphi\cdot\big(2^{\uparrow\partial_\varphi}(\kappa)\big)^{\partial_\varphi}$ leading existential quantifiers. □

C Details regarding Section 4.2

Theorem 16. There is a class of SF sentences that are Horn and Krom such that for every positive integer n the class contains a sentence φ of degree $\partial_\varphi = n$ and with a length linear in n for which any equivalent BSR sentence contains at least $\sum_{k=1}^{n} 2^{\uparrow k}(n)$ leading existential quantifiers.

Proof. Let $n \geq 1$ be some positive integer. Consider the following first-order sentence in which the sets $\{x_1, \ldots, x_n\}$ and $\{y_1, \ldots, y_n\}$ are separated:

$$\varphi := \forall x_n \exists y_n \ldots \forall x_1 \exists y_1 . \bigwedge_{i=1}^{4n} \big(P_i(x_1, \ldots, x_n) \leftrightarrow Q_i(y_1, \ldots, y_n)\big) \ .$$

Notice that we change the orientation of the indices in the quantifier prefix in this proof.

In order to construct a particular model of φ, we inductively define the following sets: $\mathcal{S}_1 := \big\{S \subseteq [4n] \;\big|\; |S| = 2n\big\}$, $\mathcal{S}_{k+1} := \big\{S \in \mathfrak{P}\mathcal{S}_k \;\big|\; |S| = \frac{1}{2} \cdot |\mathcal{S}_k|\big\}$ for every $k > 1$. Hence, we observe that

$$\begin{aligned}
|\mathcal{S}_1| &= \tbinom{4n}{2n} \geq \big(\tfrac{4n}{2n}\big)^{2n} = 2^{2n}, \\
|\mathcal{S}_2| &= \tbinom{|\mathcal{S}_1|}{|\mathcal{S}_1|/2} \geq \big(\tfrac{|\mathcal{S}_1|}{|\mathcal{S}_1|/2}\big)^{|\mathcal{S}_1|/2} = 2^{|\mathcal{S}_1|/2} \geq 2^{2^{2n}/2} = 2^{2^{2n-1}}, \\
&\vdots \\
|\mathcal{S}_n| &= \tbinom{|\mathcal{S}_{n-1}|}{|\mathcal{S}_{n-1}|/2} \geq 2^{|\mathcal{S}_{n-1}|/2} \geq 2^{2^{\vdots^{2^{2n-1}-1}}-1} \geq 2^{\uparrow n}(2n-(n-1)) = 2^{\uparrow n}(n+1),
\end{aligned}$$

where the inequality $\binom{n}{k} \geq (n/k)^k$ can be found in T. Cormen et al. *Introduction to Algorithms* (McGraw-Hill Higher Education, 2nd edition, 2001.), page 1097, for example.

Having the sets $\mathcal{S}_k$, we now define the structure $\mathcal{A}$ as follows:

- $\mathfrak{U}_\mathcal{A} := \bigcup_{k=1}^{n} \big\{\mathfrak{a}_S^{(k)}, \mathfrak{b}_S^{(k)} \;\big|\; S \in \mathcal{S}_k\big\}$,
- $P_i^\mathcal{A} := \big\{\langle \mathfrak{a}_{S_1}^{(1)}, \ldots, \mathfrak{a}_{S_n}^{(n)}\rangle \in \mathfrak{U}_\mathcal{A}^n \;\big|\; i \in S_1 \in S_2 \in \ldots \in S_n\big\}$ for $i = 1, \ldots, 4n$, and
- $Q_i^\mathcal{A} := \big\{\langle \mathfrak{b}_{S_1}^{(1)}, \ldots, \mathfrak{b}_{S_n}^{(n)}\rangle \in \mathfrak{U}_\mathcal{A}^n \;\big|\; i \in S_1 \in S_2 \in \ldots \in S_n\big\}$ for $i = 1, \ldots, 4n$.

Clearly, for any choice of $S_1, \ldots, S_n$ and every i, $1 \leq i \leq 4n$, we have

$$\mathcal{A}, \big[x_1 \mapsto \mathfrak{a}_{S_1}^{(1)}, \ldots, x_n \mapsto \mathfrak{a}_{S_n}^{(n)}, y_1 \mapsto \mathfrak{b}_{S_1}^{(1)}, \ldots, y_n \mapsto \mathfrak{b}_{S_n}^{(n)}\big] \models P_i(x_1, \ldots, x_n) \leftrightarrow Q_i(y_1, \ldots, y_n) \ .$$

For any other choice of tuples $\langle \mathfrak{c}_1, \ldots, \mathfrak{c}_n\rangle$, i.e. there do not exist sets $S_1 \in \mathcal{S}_1, \ldots, S_n \in \mathcal{S}_n$ such that $\langle \mathfrak{c}_1, \ldots, \mathfrak{c}_n\rangle$ equals $\langle \mathfrak{a}_{S_1}^{(1)}, \ldots, \mathfrak{a}_{S_n}^{(n)}\rangle$ or $\langle \mathfrak{b}_{S_1}^{(1)}, \ldots, \mathfrak{b}_{S_n}^{(n)}\rangle$, we observe $\mathcal{A}, [x_1 \mapsto \mathfrak{c}_1, \ldots, x_n \mapsto \mathfrak{c}_n] \not\models P_i(x_1, \ldots, x_n)$ and $\mathcal{A}, [y_1 \mapsto \mathfrak{c}_1, \ldots, y_n \mapsto \mathfrak{c}_n] \not\models Q_i(y_1, \ldots, y_n)$ for every i. Hence,

$$\mathcal{A}, [x_1 \mapsto \mathfrak{c}_1, \ldots, x_n \mapsto \mathfrak{c}_n, y_1 \mapsto \mathfrak{c}_1, \ldots, y_n \mapsto \mathfrak{c}_n] \models \bigwedge_{i=1}^{4n} P_i(x_1, \ldots, x_n) \leftrightarrow Q_i(y_1, \ldots, y_n) \ .$$

Consequently, $\mathcal{A}$ is a model of φ.

Consider the following simple two-player game with Players $\mathfrak{A}$ and $\mathfrak{B}$ where both players have complete and instantaneous knowledge about all moves that are made by either player. In the first round $\mathfrak{A}$ moves first by picking some domain element $\mathfrak{a}_{S_{\mathfrak{A},n}}^{(n)}$ for some set $S_{\mathfrak{A},n} \in \mathcal{S}_n$. $\mathfrak{B}$ knows about $\mathfrak{A}$'s choice and answers by picking a domain element $\mathfrak{b}_{S_{\mathfrak{B},n}}^{(n)}$ for some set $S_{\mathfrak{B},n} \in \mathcal{S}_n$. The game continues for $n-1$ more rounds, where in every round Player $\mathfrak{A}$ picks a domain element $\mathfrak{a}_{S_{\mathfrak{A},j}}^{(j)}$ with $S_{\mathfrak{A},j} \in S_{\mathfrak{A},j+1}$ and $\mathfrak{B}$ answers by picking some $\mathfrak{b}_{S_{\mathfrak{B},j}}^{(j)} \in S_{\mathfrak{B},j+1}$. Hence, in the last round the

chosen domain elements $\mathfrak{a}^{(1)}_{S_{\mathfrak{A},1}}$ and $\mathfrak{b}^{(1)}_{S_{\mathfrak{B},1}}$ are such that $S_{\mathfrak{A},1}$ and $S_{\mathfrak{B},1}$ are both nonempty subsets of $[4n]$. Player $\mathfrak{A}$ wins if and only if

$$\mathcal{A}, \left[x_1 \mapsto \mathfrak{a}^{(1)}_{S_{\mathfrak{A},1}}, \ldots, x_n \mapsto \mathfrak{a}^{(n)}_{S_{\mathfrak{A},n}}, y_1 \mapsto \mathfrak{b}^{(1)}_{S_{\mathfrak{B},1}}, \ldots, y_n \mapsto \mathfrak{b}^{(n)}_{S_{\mathfrak{B},n}}\right] \not\models P_i(x_1, \ldots, x_n) \leftrightarrow Q_i(y_1, \ldots, y_n)$$

for some $i \in [4n]$, and Player $\mathfrak{B}$ wins if and only if

$$\mathcal{A}, \left[x_1 \mapsto \mathfrak{a}^{(1)}_{S_{\mathfrak{A},1}}, \ldots, x_n \mapsto \mathfrak{a}^{(n)}_{S_{\mathfrak{A},n}}, y_1 \mapsto \mathfrak{b}^{(1)}_{S_{\mathfrak{B},1}}, \ldots, y_n \mapsto \mathfrak{b}^{(n)}_{S_{\mathfrak{B},n}}\right] \models P_i(x_1, \ldots, x_n) \leftrightarrow Q_i(y_1, \ldots, y_n)$$

for every $i \in [4n]$. Since $\mathcal{A}$ is a model of φ, there must exist a winning strategy for $\mathfrak{B}$.

Claim: There is exactly one winning strategy for $\mathfrak{B}$, namely, for every $j = n, \ldots, 1$ Player $\mathfrak{B}$ picks the element $\mathfrak{b}^{(j)}_{S_{\mathfrak{A},j}}$ in round $n - j + 1$, i.e. for every j we have $S_{\mathfrak{B},j} = S_{\mathfrak{A},j}$.

Proof: It is easy to see that the described strategy is a winning strategy for $\mathfrak{B}$.

Assume $\mathfrak{B}$ deviates from this strategy. This means there exists some j_*, $1 \leq j_* \leq n$, such that $\mathfrak{B}$ did not adhere to the described strategy in the $(n - j_* + 1)$st round, i.e. $S_{\mathfrak{B},j_*} \neq S_{\mathfrak{A},j_*}$.

We show by induction on j_* that $\mathfrak{A}$ has a winning strategy from this deviation point on.

For the base case $j = 1$ we consider two distinct nonempty sets $S_{\mathfrak{A},1}, S_{\mathfrak{B},1} \subseteq [4n]$. There must be some index i_* that belongs to one of the two sets but not to the other, i.e. $i_* \in (S_{\mathfrak{A},1} \cup S_{\mathfrak{B},1}) \setminus (S_{\mathfrak{A},1} \cap S_{\mathfrak{B},1})$.

Suppose that $i_* \in S_{\mathfrak{A},1} \setminus S_{\mathfrak{B},1}$. Hence, we can construct the chain $i_* \in S_{\mathfrak{A},1} \in \ldots \in S_{\mathfrak{A},n}$, by definition of the allowed moves. This entails $\mathcal{A}, [x_1 \mapsto \mathfrak{a}^{(1)}_{S_{\mathfrak{A},1}}, \ldots, x_n \mapsto \mathfrak{a}^{(n)}_{S_{\mathfrak{A},n}}] \models P_{i_*}(x_1, \ldots, x_n)$.

On the other hand, we know $\mathcal{A}, [y_1 \mapsto \mathfrak{b}^{(1)}_{S_{\mathfrak{B},1}}, \ldots, y_n \mapsto \mathfrak{b}^{(n)}_{S_{\mathfrak{B},n}}] \not\models Q_{i_*}(y_1, \ldots, y_n)$, because of $i_* \notin S_{\mathfrak{B},1}$. Hence, $\mathfrak{A}$ wins and the chosen strategy cannot be a winning strategy for $\mathfrak{B}$.

The case where $i_* \in S_{\mathfrak{B},1} \setminus S_{\mathfrak{A},1}$ is symmetric and $\mathfrak{A}$ also wins.

For the inductive case we fix some $j_* > 1$. Since $S_{\mathfrak{A},j_*}$ and $S_{\mathfrak{B},j_*}$ are distinct but have the same number of elements, there is some set $S' \in S_{\mathfrak{A},j_*} \setminus S_{\mathfrak{B},j_*}$. If $\mathfrak{A}$ picks $\mathfrak{a}^{(j_*-1)}_{S_{\mathfrak{A},j_*-1}} := \mathfrak{a}^{(j_*-1)}_{S'}$ in the following round, we have $S_{\mathfrak{B},j_*-1} \neq S_{\mathfrak{A},j_*-1}$ for any choice $\mathfrak{b}^{(j_*-1)}_{S_{\mathfrak{B},j_*-1}}$ that $\mathfrak{B}$ could possibly make. By induction, $\mathfrak{A}$ has a winning strategy starting from the next round of the game. Hence, there is a winning strategy starting from the current round. ◊

The just proved claim would still hold true if we allowed $\mathfrak{B}$ to freely pick any element of the domain $\mathfrak{U}_{\mathcal{A}}$ at every round. The reason is that for any choice of elements $\mathfrak{a}^{(n)}_{S_{\mathfrak{A},n}}, \ldots, \mathfrak{a}^{(1)}_{S_{\mathfrak{A},1}}$ made by $\mathfrak{A}$ with $S_{\mathfrak{A},1} \in \ldots \in S_{\mathfrak{A},n} \in \mathcal{S}_n$ we know that $S_{\mathfrak{A},1}$ is nonempty. Hence, we can always find some $i_* \in S_{\mathfrak{A},1}$ such that $\langle \mathfrak{a}^{(n)}_{S_{\mathfrak{A},n}}, \ldots, \mathfrak{a}^{(1)}_{S_{\mathfrak{A},1}} \rangle \in P^{\mathcal{A}}_{i_*}$. On the other hand, for any sequence $\mathfrak{c}_n, \ldots, \mathfrak{c}_1$ picked by $\mathfrak{B}$ that does not comply with the rules of the described game, we have $\langle \mathfrak{c}_n, \ldots, \mathfrak{c}_1 \rangle \notin Q^{\mathcal{A}}_{i_*}$.

This result proves the following observation for any of the $\mathfrak{b}^{(k)}_S$:

(∗) removing $\mathfrak{b}^{(k)}_S$ from $\mathcal{A}$'s domain and restricting the sets $Q^{\mathcal{A}}_i$ to subsets of $\left(\mathfrak{U}_{\mathcal{A}} \setminus \{\mathfrak{b}^{(k)}_S\}\right)^n$ yields a structure that does not satisfy φ.

The reason is simply that in this case player $\mathfrak{A}$ can always prevent $\mathfrak{B}$ from reaching a state of the game where $\mathfrak{B}$ can apply the described winning strategy.

We have already analyzed the size of the sets $\mathcal{S}_k$. Due to the observed lower bounds, we know that $\mathfrak{U}_{\mathcal{A}}$ contains at least $\sum_{k=1}^{n} 2^{\uparrow k}(n)$ elements of the form $\mathfrak{b}^{(k)}_S$.

Next, we argue that any sentence φ_* (in prenex normal form) that is semantically equivalent to φ and starts with a quantifier prefix of the form $\exists^*\forall^*$ contains at least $\sum_{k=1}^{n} 2^{\uparrow k}(n)$ leading existential quantifiers.

Let $\varphi_* := \exists y_1 \ldots y_m \forall x_1 \ldots x_\ell . \chi_*$ (with χ_* being quantifier-free) be a sentence with minimal m that is semantically equivalent to φ. Since $\mathcal{A}$ is also a model of φ_*, we know that there is a sequence

of elements $\mathfrak{c}_1, \ldots, \mathfrak{c}_m$ taken from the domain $\mathfrak{U}_\mathcal{A}$ such that $\mathcal{A}, [y_1 \mapsto \mathfrak{c}_1, \ldots, y_m \mapsto \mathfrak{c}_m] \models \forall x_1 \ldots x_\ell . \chi_*$. Consequently, we can extend $\mathcal{A}$ to a model $\mathcal{A}_*$ (over the same domain) of the Skolemized formula $\varphi_{\mathrm{Sk}} := \forall x_1 \ldots x_\ell . \chi_*\big[y_1/c_1, \ldots, y_m/c_m\big]$ by adding $c_j^{\mathcal{A}_*} := \mathfrak{c}_j$ for $j = 1, \ldots, m$. On the other hand, every model of the Skolemized formula φ_{Sk} immediately yields a model of φ_*.

The signature underlying φ_{Sk} comprises exactly the constant symbols $c_1, \ldots, c_m$ and does not contain any other function symbols. Suppose $m < \sum_{k=1}^{n} 2^{\uparrow k}(n)$. Hence, there is some $\mathfrak{b}_S^{(k)}$ with $S \in \mathcal{S}_k$ such that for every j we have $c_j^{\mathcal{A}_*} \neq \mathfrak{b}_S^{(k)}$. By the Substructure Lemma, the following substructure $\mathcal{B}$ of $\mathcal{A}_*$ constitutes a model of φ_{Sk}: $\mathfrak{U}_\mathcal{B} := \mathfrak{U}_{\mathcal{A}_*} \setminus \{\mathfrak{b}_S^{(k)}\}$, $P_i^\mathcal{B} := P_i^{\mathcal{A}_*} \cap \mathfrak{U}_\mathcal{B}^n = P_i^{\mathcal{A}_*}$ and $Q_i^\mathcal{B} := Q_i^{\mathcal{A}_*} \cap \mathfrak{U}_\mathcal{B}^n$ for every i, and $c_j^\mathcal{B} := c_j^{\mathcal{A}_*}$ for every j.

But then $\mathcal{B}$ must also be a model of both φ_* and φ, since every model of φ_{Sk} is a model of φ_*, and because we assumed φ_* and φ to be semantically equivalent. This contradicts Observation $(*)$, and thus we must have $m \geq \sum_{k=1}^{n} 2^{\uparrow k}(n)$.

Since every atom $Q_i(y_1, \ldots, y_n)$ contains n variables from existential quantifier blocks that are interspersed with universal quantifier blocks, the degree ∂_φ of φ is n. Moreover, φ can easily be transformed into a CNF that is Horn and Krom at the same time. Hence, the theorem holds. □

D Details regarding Section 5

The following theorems illustrate why the domino problems defined in Section 5 readily lend themselves to deriving lower bounds on worst-case computing time.

Theorem 28 ([3], Theorem 6.1.2)**.** Let M be a *simple* nondeterministic one-tape Turing machine with input alphabet Γ. Then there is a domino system $\mathfrak{D} = \langle \mathcal{D}, \mathcal{H}, \mathcal{V} \rangle$ and a linear-time reduction which takes any input $w \in \Gamma^*$ to a word $\overline{D} \in \mathcal{D}^*$ with $|w| = |\overline{D}|$ such that

- If M accepts w in time t_0 with space s_0, then $\mathfrak{D}$ tiles $\mathbb{Z}_s \times \mathbb{Z}_t$ with initial condition $\overline{D}$ for all $s \geq s_0 + 2$ and $t \geq t_0 + 2$;
- If M does not accept w, then $\mathfrak{D}$ does not tile $\mathbb{Z}_s \times \mathbb{Z}_t$ with initial condition $\overline{D}$ for any $s, t \geq 2$.

By a *simple* Turing machine the authors of [3] mean a nondeterministic one-tape Turing machine M over the input alphabet Γ that meets the following conditions:

> "The alphabet of M contains Γ and at least one other symbol $\square$ (blank). M works on a semi-infinite tape and never tries to move left from the left-most tape cell. At every stage of the computation there is some s such the tape cells $0, \ldots, s$ contain only non-blank symbols, all other tape cells contain $\square$; in particular, to the right of a blank only other blanks may appear. Furthermore, we assume that M has a unique accepting configuration: the machine is in the unique accepting state q_a, the tape contains only blanks and the head is in position 0.
>
> These conditions do not restrict computational power. Every language accepted in time $T(n)$ and space $S(n)$ by some one-tape nondeterministic Turing machine is accepted within the same time and space bounds by a simple Turing machine, as long as $S(n), T(n) \geq 2n$." [3], page 243

Theorem 29 ([3], Theorem 6.1.6)**.** We call a function $T : \mathbb{N} \to \mathbb{N}$ *time constructible* if there exists a deterministic Turing machine making precisely $T(n)$ steps on inputs of length n.

Let $T : \mathbb{N} \to \mathbb{N}$ be a time-constructible function with $T(c'n)^2 \in o(T(n))$ for some constant $c' > 0$. There exists a domino system $\mathfrak{D}$ and a constant $c > 0$ such that $\mathrm{DOMINO}(\mathfrak{D}, T(n)) \notin \mathrm{NTIME}(T(cn))$.

Recall our notation for the *tetration operation* $2^{\uparrow k}(m)$, which we defined inductively: $2^{\uparrow 0}(m) := m$ and $2^{\uparrow k+1}(m) := 2^{\left(2^{\uparrow k}(m)\right)}$. In addition, we shall use the short-hand $2^{\uparrow k}$ to abbreviate $2^{\uparrow k}(2)$.

Since we intend to plug in the function $2^{\uparrow n}$ for $T(n)$, we need to verify the condition $\left(2^{\uparrow c'n}\right)^2 \in o(2^{\uparrow n})$ for some positive constant c'. Setting $c' := \frac{1}{2}$ entails the following.

Lemma 30. For every positive constant $c > 0$ there exists some positive integer $n_0 > 0$ such that for every $n \geq n_0$ we have

$$\left(2^{\uparrow\lceil n/2\rceil}\right)^2 \leq c \cdot 2^{\uparrow n} \ .$$

Proof. We distinguish two cases:

Suppose $c \geq 1$. We observe

(1) $\left(2^{\uparrow\lceil n/2\rceil}\right)^2 \ = \ 2^{2\cdot 2^{2^{\uparrow\lceil n/2\rceil-2}}} \ = \ 2^{2^{1+2^{\uparrow\lceil n/2\rceil-2}}}$, and

(2) $c \cdot 2^{\uparrow n} \ \geq \ 2^{\uparrow n} \ = \ 2^{2^{2^{\uparrow n-2}}}$.

Hence, it suffices to show there is some $n_0 \geq 2$ such that $1 + 2^{\uparrow\lceil n/2\rceil-2} \ \leq \ 2^{\uparrow n-2}$ holds for every $n \geq n_0$. One possible choice is $n_0 = 4$

Suppose $0 < c < 1$. We set $d := \frac{1}{c}$. Due to $d > 1$, $\log_2 d$ is defined. Moreover, we observe

$$\begin{aligned} d \cdot \left(2^{\uparrow\lceil n/2\rceil}\right)^2 \ &= \ 2^{\log_2 d} \cdot \left(2^{\uparrow\lceil n/2\rceil}\right)^2 \\ &= \ 2^{\log_2 d + 2\cdot 2^{\uparrow\lceil n/2\rceil-1}} \\ &\leq \ 2^{\log_2 d \cdot 2^{\uparrow\lceil n/2\rceil-1} + 2\cdot 2^{\uparrow\lceil n/2\rceil-1}} \\ &= \ 2^{(\log_2 d+2)\cdot 2^{\uparrow\lceil n/2\rceil-1}} \\ &= \ 2^{2^{\log_2(\log_2 d+2)}\cdot 2^{\uparrow\lceil n/2\rceil-2}} \\ &= \ 2^{2^{\log_2(\log_2 d+2)+2^{\uparrow\lceil n/2\rceil-2}}} \ . \end{aligned}$$

Hence, in order to prove that there is some n_0 such that $d \cdot \left(2^{\uparrow\lceil n/2\rceil}\right)^2 \leq 2^{\uparrow n}$ holds for every $n \geq n_0$, it suffices to show that there is some n_0 such that

$$d' + 2^{\uparrow\lceil n/2\rceil-2} \ \leq \ 2^{\uparrow n-2}$$

where $d' := \log_2(\log_2 d + 2)$. The proof thus boils down to asking whether the difference $2^{\uparrow n-2} - 2^{\uparrow\lceil n/2\rceil-2}$ exceeds any constant value d' for sufficiently large n. This is certainly the case. □

Lemma 31. Let $k \geq 1$ be some positive integer. For every positive constant $c > 0$ there exists some positive integer $n_0 \geq 1$ such that for every $n \geq n_0$ we have

$$\left(2^{\uparrow k}(\lceil n/4\rceil)\right)^2 \leq c \cdot 2^{\uparrow k}(n) \ .$$

Proof. Let $k = 1$. We observe

$$\left(2^{\uparrow 1}(n/4)\right)^2 = 2^{2\cdot n/4} = 2^{n/2} \ .$$

If $c \geq 1$, then $2^{n/2} \leq c \cdot 2^n$ is obvious.

Assume $0 < c < 1$ and set $d := \frac{1}{c}$. Hence, $d > 1$. If remains to show $d \cdot 2^{n/2} \leq 2^n$ for every sufficiently large n. Due to $2^n = 2^{n/2+n/2} = 2^{n/2} \cdot 2^{n/2}$, we observe $d \cdot 2^{n/2} \leq 2^{n/2} \cdot 2^{n/2}$ if and only if $d \leq 2^{n/2}$. But the latter certainly holds for sufficiently large n.

Let $k = 2$. If $c \geq 1$, then $2^{\uparrow 2}(n) \leq c \cdot 2^{\uparrow 2}(n)$. It thus suffices to show $2^{2^{n/4}} \leq 2^{2^n}$ or, equivalently, $\frac{1}{4}n + 1 \leq n$ for sufficiently large n. But this is obviously true.

Assume $0 < c < 1$ and define $d := \frac{1}{c}$. It follows that $d > 1$. It remains to show $d \cdot 2^{2^{n/4+1}} \leq 2^{2^n}$ for sufficiently large n. Due to $2^{2^n} = 2^{2^{n/2}\cdot 2^{n/2}} \geq 2^{2^{n/2}+2^{n/2}} = 2^{2^{n/2}} \cdot 2^{2^{n/2}}$ with $d \leq 2^{2^{n/2}}$ and $2^{2^{n/4+1}} \leq 2^{2^{n/2}}$ for sufficiently large n, we also observe $d \cdot 2^{2^{n/4+1}} \leq 2^{2^{n/2}} \cdot 2^{2^{n/2}} \leq 2^{2^n}$ for sufficiently large n.

Analogous arguments holds for every $k \geq 2$. □

E Details regarding Section 5.1

We investigate the properties of models of the SF sentences given in Section 5.1. For the remainder of this subsection we therefore consider a model $\mathcal{A}$ of the sentence $\psi_1 \wedge \ldots \wedge \psi_{16} \wedge \chi_1 \wedge \chi_2 \wedge \chi_3$.

Definition 32. We define the following sets and relations:

- $\mathcal{I}_\ell := \{\mathfrak{a} \in \mathfrak{U}_\mathcal{A} \mid \mathcal{A}, [j \mapsto \mathfrak{a}] \models L(\underline{\ell}, j)\}$ for every $\ell = 0, \ldots, \kappa$.
- $\prec_\ell \subseteq \mathcal{I}_\ell \times \mathcal{I}_\ell$ for every $\ell = 0, \ldots, \kappa$ such that $\mathfrak{a} \prec_\ell \mathfrak{a}'$ holds if and only if $\mathcal{A}, [j \mapsto \mathfrak{a}, j' \mapsto \mathfrak{a}'] \models \mathrm{Succ}(\underline{\ell}, j, j')$.
- $\mathcal{F}_\ell := \{f : \mathcal{I}_\ell \to \{0,1\} \mid f$ is total and for every $\mathfrak{a}$ for which $\mathcal{A}, [j \mapsto \mathfrak{a}] \models \mathrm{MaxIdx}(\underline{\ell}, j)$ we have that $f(\mathfrak{a}) = 1$ entails $f(\mathfrak{b}) = 0$ for every $\mathfrak{b} \neq \mathfrak{a}\}$
  for every $\ell = 0, \ldots, \kappa - 1$.
- For every $\ell = 0, \ldots, \kappa - 1$ and two total functions $f, g : \mathcal{I}_\ell \to \{0,1\}$ we write $f \sqsubset_\ell g$ if and only if
  - there is some integer p and a unique chain $\mathfrak{a}_1 \prec_\ell \ldots \prec_\ell \mathfrak{a}_p$ comprising all elements in $\mathcal{I}_\ell$, and
  - incrementing the bit string $f(\mathfrak{a}_1)f(\mathfrak{a}_2)\ldots f(\mathfrak{a}_p)$ (interpreted as number encoded in binary where the leftmost bit is the least significant one) by one yields $g(\mathfrak{a}_1)g(\mathfrak{a}_2)\ldots g(\mathfrak{a}_p)$.
- $\mathcal{S}_{\ell,f} := \{\mathfrak{a} \in \mathcal{I}_\ell \mid \mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto \mathfrak{b}] \models J(\underline{\ell}, j, i, f(\mathfrak{b}))$ for every $\mathfrak{b} \in \mathcal{I}_{\ell-1}\}$
  for every $\ell = 1, \ldots, \kappa$ and every $f \in \mathcal{F}_{\ell-1}$.
- For every $\ell = 1, \ldots, \kappa$ and two elements $\mathfrak{a}, \mathfrak{a}' \in \mathcal{I}_\ell$ we write $\mathfrak{a} \sim_\ell \mathfrak{a}'$ if and only if for every $\mathfrak{b} \in \mathcal{I}_{\ell-1}$ we observe
  - $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto \mathfrak{b}] \models J(\underline{\ell}, j, i, 0)$ if and only if $\mathcal{A}, [j' \mapsto \mathfrak{a}', i \mapsto \mathfrak{b}] \models J(\underline{\ell}, j', i, 0)$ and
  - $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto \mathfrak{b}] \models J(\underline{\ell}, j, i, 1)$ if and only if $\mathcal{A}, [j' \mapsto \mathfrak{a}', i \mapsto \mathfrak{b}] \models J(\underline{\ell}, j', i, 1)$.

Proposition 33. For every ℓ, $1 \leq \ell \leq \kappa$, $\sim_\ell$ is an equivalence relation.

Proposition 34. For all distinct ℓ, ℓ', $0 \leq \ell, \ell' \leq \kappa$, $\mathcal{I}_\ell$ and $\mathcal{I}_{\ell'}$ are disjoint.

Lemma 35. Let $\mathfrak{a}, \mathfrak{a}' \in \mathcal{I}_\ell$ for some ℓ, $1 \leq \ell \leq \kappa$. Let $f, g : \mathcal{I}_{\ell-1} \to \{0,1\}$ be two total functions such that for every $\mathfrak{b} \in \mathcal{I}_{\ell-1}$ we have $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto \mathfrak{b}] \models J(\underline{\ell}, j, i, f(\mathfrak{b}))$ and $\mathcal{A}, [j' \mapsto \mathfrak{a}', i \mapsto \mathfrak{b}] \models J(\underline{\ell}, j', i, g(\mathfrak{b}))$. Moreover, assume there is a unique chain $\mathfrak{c}_1 \prec_{\ell-1} \ldots \prec_{\ell-1} \mathfrak{c}_p$ comprising all elements in $\mathcal{I}_{\ell-1}$.

Then $\mathfrak{a} \prec_\ell \mathfrak{a}'$ implies $f \sqsubset_{\ell-1} g$.

Proof. Since $\mathfrak{c}_1$ is the only element in $\mathcal{I}_{\ell-1}$ for which there is no element $\mathfrak{c}' \in \mathcal{I}_{\ell-1}$ with $\mathfrak{c}' \prec_{\ell-1} \mathfrak{c}_1$, $\mathcal{A} \models \psi_2 \wedge \psi_3$ implies $\mathcal{A}, [i \mapsto \mathfrak{c}_1] \models \mathrm{MinIdx}(\underline{\ell - 1}, i)$.

Let $f^* : \mathcal{I}_{\ell-1} \to \{0,1\}$ be defined such that

- $f^*(\mathfrak{b}) = 0$ if and only if $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto \mathfrak{b}] \models J^*(\underline{\ell}, j, i, 0)$ and
- $f^*(\mathfrak{b}) = 1$ if and only if $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto \mathfrak{b}] \models J^*(\underline{\ell}, j, i, 1)$.

This function is well-defined because of $\mathcal{A} \models \psi_{13}$. Due to $\mathcal{A} \models \chi_2 \wedge \psi_{14} \wedge \psi_{15}$ it is also total.

$\mathcal{A} \models \psi_{14}$ enforces $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto \mathfrak{c}_1] \models J^*(\underline{\ell}, j, i, 1)$, i.e. $f^*(\mathfrak{c}_1) = 1$. Moreover, for any k with $1 < k \leq p$, we have $f^*(\mathfrak{c}_k) = 1$ if and only if $f(\mathfrak{c}_1) = \ldots = f(\mathfrak{c}_{k-1}) = 1$, because of $\mathcal{A} \models \psi_{15}$.

$\mathcal{A} \models \psi_9$ together with our assumption $\mathfrak{a} \prec_\ell \mathfrak{a}'$ translates to the following property, which we phrase in terms of operations on bits:

$(**)$ for every k, $1 \leq k \leq p$, we observe $g(\mathfrak{c}_k) = f(\mathfrak{c}_k) \oplus f^*(\mathfrak{c}_k)$

where $\oplus$ denotes exclusive or. But this corresponds to an increase of the bit string $f(\mathfrak{c}_1)\dots f(\mathfrak{c}_p)$ by one (where $f(\mathfrak{c}_1)$ is the least significant bit).

Hence, $f \sqsubset_{\ell-1} g$. □

Lemma 36. Let $\mathfrak{a}, \mathfrak{a}' \in \mathcal{I}_\ell$ for some ℓ, $1 \leq \ell \leq \kappa$.

If $\mathfrak{a}$ and $\mathfrak{a}'$ belong to the same $\mathcal{S}_{\ell,f}$ for some function $f \in \mathcal{F}_{\ell-1}$, then $\mathfrak{a} \sim_\ell \mathfrak{a}'$.

Proof. By totality of f, it follows that for every $\mathfrak{b} \in \mathcal{I}_{\ell-1}$ we have $\mathcal{A}, [j{\mapsto}\mathfrak{a}, i{\mapsto}\mathfrak{b}] \models J(\underline{\ell}, j, i, \underline{f(\mathfrak{b})})$ and $\mathcal{A}, [j'{\mapsto}\mathfrak{a}', i{\mapsto}\mathfrak{b}] \models J(\underline{\ell}, j', i, \underline{f(\mathfrak{b})})$. Moreover, $\mathcal{A} \models \psi_{13}$ entails $\mathcal{A}, [j{\mapsto}\mathfrak{a}, i{\mapsto}\mathfrak{b}] \not\models J(\underline{\ell}, j, i, \overline{f(\mathfrak{b})})$ and $\mathcal{A}, [j'{\mapsto}\mathfrak{a}', i{\mapsto}\mathfrak{b}] \not\models J(\underline{\ell}, j', i, \overline{f(\mathfrak{b})})$ where $\overline{f(\mathfrak{b})} := 1 - f(\mathfrak{b})$ is the complement of $f(\mathfrak{b})$. This results in $\mathfrak{a} \sim_\ell \mathfrak{a}'$. □

Lemma 37. Let $\mathfrak{a}, \mathfrak{a}' \in \mathcal{I}_\ell$ for some ℓ, $1 \leq \ell \leq \kappa$. Moreover, let f be some function in $\mathcal{F}_{\ell-1}$.

If $\mathfrak{a}$ belongs to $\mathcal{S}_{\ell,f}$ and we have $\mathfrak{a} \sim_\ell \mathfrak{a}'$, then $\mathfrak{a}' \in \mathcal{S}_{\ell,f}$.

Proof. By totality of f, it follows that for every $\mathfrak{b} \in \mathcal{I}_{\ell-1}$ we have $\mathcal{A}, [j{\mapsto}\mathfrak{a}, i{\mapsto}\mathfrak{b}] \models J(\underline{\ell}, j, i, f(\mathfrak{b}))$. Moreover, $\mathcal{A} \models \psi_{13}$ entails $\mathcal{A}, [j{\mapsto}\mathfrak{a}, i{\mapsto}\mathfrak{b}] \not\models J(\underline{\ell}, j, i, \overline{f(\mathfrak{b})})$ where $\overline{f(\mathfrak{b})} := 1 - f(\mathfrak{b})$ is the complement of $f(\mathfrak{b})$.

Since we assume $\mathfrak{a} \sim_\ell \mathfrak{a}'$, we know that these properties transfer to $\mathfrak{a}'$. Hence, for every $\mathfrak{b} \in \mathcal{I}_{\ell-1}$ we observe $\mathcal{A}, [j'{\mapsto}\mathfrak{a}', i{\mapsto}\mathfrak{b}] \models J(\underline{\ell}, j', i, f(\mathfrak{b}))$ and $\mathcal{A}, [j'{\mapsto}\mathfrak{a}', i{\mapsto}\mathfrak{b}] \not\models J(\underline{\ell}, j', i, \overline{f(\mathfrak{b})})$. Consequently, $\mathfrak{a}' \in \mathcal{S}_{\ell,f}$. □

Lemma 38. Consider some ℓ, $0 \leq \ell \leq \kappa - 1$. If there is a unique chain $\mathfrak{a}_1 \prec_\ell \dots \prec_\ell \mathfrak{a}_p$ comprising all elements in $\mathcal{I}_\ell$, then we can uniquely arrange all functions in $\mathcal{F}_\ell$ into a chain $f_1 \sqsubset_\ell \dots \sqsubset_\ell f_{p'}$ where $p' = |\mathcal{F}_\ell| = 2^{p-1} + 1$.

Proof. Let $\{0,1\}^p$ be the set of all bit strings of length p. If we interpret each of them as a number encoded in binary (where we assume the leftmost bit to be the least significant one), we can uniquely arrange the $2^{p-1}+1$ smallest bit strings in $\{0,1\}^p$ into a chain

$$\overline{b}_0 < \overline{b}_1 < \dots < \overline{b}_{2^{p-1}-1} < \overline{b}_{2^{p-1}}$$

where the indices reflect the represented numerical value and $<$ is intended to be the usual ordering based on this value. Since we assume the rightmost bit to be the most significant one, it is 0 in $\overline{b}_0, \dots, \overline{b}_{2^{p-1}-1}$. Accordingly, $\overline{b}_{2^{p-1}}$ is the bit string with all zeros except for the most significant bit, i.e. $\overline{b}_{2^{p-1}} = 0\dots01$.

Obviously, the following mapping μ induces a one-to-one correspondence between bit strings and all the mappings in $\mathcal{F}_\ell$: $\mu(f) := f(\mathfrak{a}_1)f(\mathfrak{a}_2)\dots f(\mathfrak{a}_p)$. By definition of $\sqsubset_\ell$, we have $f \sqsubset_\ell g$ if and only if $\mu(f) + 1 = \mu(g)$.

Consequently, we obtain the chain $\mu^{-1}(\overline{b}_0) \sqsubset_\ell \mu^{-1}(\overline{b}_1) \sqsubset_\ell \dots \sqsubset_\ell \mu^{-1}(\overline{b}_{2^{p-1}-1}) \sqsubset_\ell \mu^{-1}(\overline{b}_{2^{p-1}})$. □

Lemma 39. We observe that $\mathcal{I}_0 = \{c_1^{\mathcal{A}}, \dots, c_\mu^{\mathcal{A}}\}$ with $|\mathcal{I}_0| = \mu$. Moreover, there is a unique chain $c_1^{\mathcal{A}} \prec_0 \dots \prec_0 c_\mu^{\mathcal{A}}$.

Proof. $\mathcal{A} \models \chi_1$ entails that $\mathcal{I}_0 \subseteq \{c_1^{\mathcal{A}}, \dots, c_\mu^{\mathcal{A}}\}$. Due to $\mathcal{A} \models \psi_8$ we have $c_1^{\mathcal{A}} \prec_0 \dots \prec_0 c_\mu^{\mathcal{A}}$. By $\mathcal{A} \models \psi_6$, $\mathfrak{a} \prec_0 \mathfrak{b}$ entails $\mathfrak{a}, \mathfrak{b} \in \mathcal{I}_0$. Hence, $\{c_1^{\mathcal{A}}, \dots, c_\mu^{\mathcal{A}}\} \subseteq \mathcal{I}_0$

We next show that all $c_1^{\mathcal{A}}, \dots, c_\mu^{\mathcal{A}}$ are pairwise distinct.

Claim: For every index $j \geq 2$ the first j elements $c_1^{\mathcal{A}}, \dots, c_j^{\mathcal{A}}$ are distinct.

Proof: We proceed by induction on j.

For $j = 2$, $c_1^{\mathcal{A}} \neq c_2^{\mathcal{A}}$ must hold, for otherwise $c_1^{\mathcal{A}} \prec_0 c_2^{\mathcal{A}}$ contradicts $\mathcal{A} \models \psi_7$ which entails $\mathcal{A} \models \forall j.\, \neg\mathrm{Succ}(\underline{0}, j, j)$.

Let $j \geq 3$ and assume, by induction, that the elements $c_1^{\mathcal{A}}, \dots, c_{j-1}^{\mathcal{A}}$ are all pairwise distinct. Suppose there is some index i_*, $1 \leq i_* < j$, such that $c_{i_*}^{\mathcal{A}} = c_j^{\mathcal{A}}$. We distinguish two cases: In case $i_* = 1$, we have $c_{j-1}^{\mathcal{A}} \prec_0 c_1^{\mathcal{A}}$. But this contradicts $\mathcal{A} \models \psi_2$.

In case of $i_* > 1$, $\mathcal{A} \models \psi_7$ entails $c^{\mathcal{A}}_{i_*-1} = c^{\mathcal{A}}_{j-1}$, since we have $c^{\mathcal{A}}_{i_*-1} \prec_0 c^{\mathcal{A}}_{i_*}$ and $c^{\mathcal{A}}_{j-1} \prec_0 c^{\mathcal{A}}_{j}$, and since we assumed $c^{\mathcal{A}}_{i_*} = c^{\mathcal{A}}_{j}$. But this contradicts our inductive hypothesis, because $i_* - 1$ and $j - 1$ are distinct indices and thus the inductive hypothesis implies that $c^{\mathcal{A}}_{i_*-1}$ and $c^{\mathcal{A}}_{j-1}$ are distinct. ◊

The above claim entails $|\mathcal{I}_0| = \mu$.

$\mathcal{A} \models \psi_8$ entails $c^{\mathcal{A}}_1 \prec_0 \ldots \prec_0 c^{\mathcal{A}}_\mu$. By the above arguments, we know that this chain comprises all elements in $\mathcal{I}_0$. Moreover, due to $\mathcal{A} \models \psi_7$, this chain is the only chain satisfying the desired properties. □

Lemma 40. For every $\ell = 1, \ldots, \kappa$ the following properties are satisfied:

(i) For all $\mathfrak{a}, \widetilde{\mathfrak{a}} \in \mathcal{I}_\ell$ for which $\mathcal{A}, [j \mapsto \mathfrak{a}, \widetilde{j} \mapsto \widetilde{\mathfrak{a}}] \models \mathrm{eq}^\ell_{j,\widetilde{j}}$ it follows $\mathfrak{a} \sim_\ell \widetilde{\mathfrak{a}}$.

(ii) All $\mathfrak{a}, \widetilde{\mathfrak{a}} \in \mathcal{I}_\ell$ with $\mathfrak{a} \sim_\ell \widetilde{\mathfrak{a}}$ satisfy $\mathcal{A}, [j \mapsto \mathfrak{a}, \widetilde{j} \mapsto \widetilde{\mathfrak{a}}] \models \mathrm{eq}^\ell_{j,\widetilde{j}}$.

(iii) For every $f \in \mathcal{F}_{\ell-1}$ the set $\mathcal{S}_{\ell,f}$ is nonempty.

(iv) $\mathcal{I}_\ell = \bigcup_{f \in \mathcal{F}_{\ell-1}} \mathcal{S}_{\ell,f}$.

(v) For every $f \in \mathcal{F}_{\ell-1}$ the set $\mathcal{S}_{\ell,f}$ contains exactly one element.

(vi) There is a unique chain $\mathfrak{a}_1 \prec_\ell \ldots \prec_\ell \mathfrak{a}_p$ comprising all elements in $\mathcal{I}_\ell$, and $\mathcal{A}, [j \mapsto \mathfrak{a}_1] \models \mathrm{MinIdx}(\underline{\ell}, j)$ and $\mathcal{A}, [j' \mapsto \mathfrak{a}_p] \models \mathrm{MaxIdx}(\underline{\ell}, j')$ hold.

Proof. We proceed by induction on ℓ.

Base case: $\ell = 1$.

Ad (i). Due to the assumption $\mathcal{A}, [j \mapsto \mathfrak{a}, \widetilde{j} \mapsto \widetilde{\mathfrak{a}}] \models \mathrm{eq}^1_{j,\widetilde{j}}$, the construction of $\mathrm{eq}^1_{j,\widetilde{j}}$ entails

$$\mathcal{A}, [j \mapsto \mathfrak{a}, \widetilde{j} \mapsto \widetilde{\mathfrak{a}}] \models \bigwedge_{i=1}^{\mu} \Big(\big(J(\underline{1}, j, c_i, 0) \leftrightarrow J(\underline{1}, \widetilde{j}, c_i, 0) \big) \wedge \big(J(\underline{1}, j, c_i, 1) \leftrightarrow J(\underline{1}, \widetilde{j}, c_i, 1) \big) \Big) .$$

By Lemma 39, we know that $\mathfrak{b} \in \mathcal{I}_0$ entails $\mathfrak{b} = c^{\mathcal{A}}_i$ for some i, $1 \leq i \leq \mu$. Consequently, for every $\mathfrak{b} \in \mathcal{I}_0$, $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto \mathfrak{b}] \models J(\underline{1}, j, i, 0)$ holds if and only if $\mathcal{A}, [\widetilde{j} \mapsto \widetilde{\mathfrak{a}}, i \mapsto \mathfrak{b}] \models J(\underline{1}, \widetilde{j}, i, 0)$ does. This entails $\mathfrak{a} \sim_\ell \widetilde{\mathfrak{a}}$, because of $\mathcal{A} \models \psi_{13} \wedge \chi_2$.

Ad (ii). By definition of $\sim_1$, $\mathfrak{a} \sim_1 \widetilde{\mathfrak{a}}$ entails that for every $\mathfrak{b} \in \mathcal{I}_0$ we have

- $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto \mathfrak{b}] \models J(\underline{1}, j, i, 0)$ if and only if $\mathcal{A}, [\widetilde{j} \mapsto \widetilde{\mathfrak{a}}, i \mapsto \mathfrak{b}] \models J(\underline{1}, \widetilde{j}, i, 0)$ and
- $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto \mathfrak{b}] \models J(\underline{1}, j, i, 1)$ if and only if $\mathcal{A}, [\widetilde{j} \mapsto \widetilde{\mathfrak{a}}, i \mapsto \mathfrak{b}] \models J(\underline{1}, \widetilde{j}, i, 1)$.

Moreover, Lemma 39 states that $\mathcal{I}_0 = \{c^{\mathcal{A}}_1, \ldots, c^{\mathcal{A}}_\mu\}$. Since $\mathfrak{a}, \widetilde{\mathfrak{a}}$ belong to $\mathcal{I}_1$, we conclude $\mathcal{A}, [j \mapsto \mathfrak{a}, \widetilde{j} \mapsto \widetilde{\mathfrak{a}}] \models \mathrm{eq}^1_{j,\widetilde{j}}$.

Ad (iii). $\mathcal{A} \models \psi_8$ entails $c^{\mathcal{A}}_1 \prec_0 \ldots \prec_0 c^{\mathcal{A}}_\mu$. By Lemma 39, we know that this chain comprises all elements in $\mathcal{I}_0$. Moreover, due to $\mathcal{A} \models \psi_7$, this chain is the only chain satisfying this property. Hence, by Lemma 38, we can arrange all functions in $\mathcal{F}_0$ into a sequence $f_1 \sqsubset_0 \ldots \sqsubset_0 f_p$ for $p = 2^{\mu-1} + 1$. Clearly, f_1 maps every element $\mathfrak{b} \in \mathcal{I}_0$ to $f_1(\mathfrak{b}) = 0$, and f_p maps every element $\mathfrak{b} \in \mathcal{I}_0 \setminus \{c^{\mathcal{A}}_\mu\}$ to $f_p(\mathfrak{b}) = 0$ and $c^{\mathcal{A}}_\mu$ to $f_p(c^{\mathcal{A}}_\mu) = 1$. By $\mathcal{A} \models \psi_3 \wedge \psi_{10}$, we know that $\mathcal{A}, [j \mapsto d^{\mathcal{A}}_1, i \mapsto \mathfrak{b}] \models J(\underline{1}, j, i, 0)$ for every $\mathfrak{b} \in \mathcal{I}_0$. Hence, $d^{\mathcal{A}}_1 \in S_{1,f_1}$.

We next show that for every k, $1 \leq k < p$, if S_{1,f_k} is nonempty, then $S_{1,f_{k+1}}$ is nonempty. Let $\mathfrak{a}$ be an element of S_{1,f_k}. Because of $k < p$, we know that $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto c^{\mathcal{A}}_\mu] \models J(\underline{1}, j, i, 0)$. By virtue of (i) and due to $\mathcal{A} \models \psi_{16}$ we conclude that there are elements $\widetilde{\mathfrak{a}}, \widetilde{\mathfrak{a}}' \in \mathcal{I}_1$ such that $\mathfrak{a} \sim_1 \widetilde{\mathfrak{a}}$ and $\widetilde{\mathfrak{a}} \prec_1 \widetilde{\mathfrak{a}}'$. By Lemma 35, this results in $\widetilde{\mathfrak{a}}' \in \mathcal{S}_{\ell,f_{k+1}}$.

Ad (iv). As a consequence of $\mathcal{A} \models \psi_{13}$ there is a unique partial function $g_{\mathfrak{a}} : \mathcal{I}_0 \to \{0,1\}$ for every $\mathfrak{a} \in \mathcal{I}_1$ such that for every $\mathfrak{b} \in \mathcal{I}_0$ we have $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto \mathfrak{b}] \models J(\underline{1}, j, i, g_{\mathfrak{a}}(\mathfrak{b}))$ if and only if $g_{\mathfrak{a}}$ is defined for $\mathfrak{b}$. Because of $\mathcal{A} \models \chi_2$, we know that $g_{\mathfrak{a}}$ must be total.

For every $\mathfrak{a} \in \mathcal{I}_1$ where $g_{\mathfrak{a}}(c_\mu^{\mathcal{A}}) = 0$ we have $g_{\mathfrak{a}} \in \mathcal{F}_0$ and thus also $\mathfrak{a} \in \mathcal{S}_{1,g_{\mathfrak{a}}}$.

Because of $\mathcal{A} \models \psi_5 \wedge \psi_{11} \wedge \psi_{12}$ we know that $\mathcal{A} \models \mathrm{MaxIdx}(\underline{1}, e_1)$ and that $e_1^{\mathcal{A}}$ is the only element in $\mathcal{I}_1$ for which $\mathcal{A}, [j \mapsto e_1^{\mathcal{A}}, i \mapsto c_\mu^{\mathcal{A}}] \models J(\underline{1}, j, i, 1)$.

It remains to show that for every $\mathfrak{b} \in \mathcal{I}_0$ with $\mathfrak{b} \neq c_\mu^{\mathcal{A}}$ we have $\mathcal{A}, [j \mapsto e_1^{\mathcal{A}}, i \mapsto \mathfrak{b}] \models J(\underline{1}, j, i, 0)$. But this is a consequence of (iii) and the fact that the total function f_* mapping all elements $\mathfrak{b}$ in $\mathcal{I}_0$ but $c_\mu^{\mathcal{A}}$ to $f_*(\mathfrak{b}) = 0$ belongs to $\mathcal{F}_0$, and thus $\mathcal{S}_{1,f_*}$ is nonempty. In particular, $\mathcal{S}_{1,f_*} = \{e_1^{\mathcal{A}}\}$.

Consequently, $\mathcal{I}_1$ cannot contain any elements that do not lie in $\bigcup_{f \in \mathcal{F}_0} \mathcal{S}_{1,f}$.

Ad (v). Consider any set $\mathcal{S}_{1,f}$. By virtue of (iii), $\mathcal{S}_{1,f}$ contains at least one element. Suppose we are given two elements $\mathfrak{a}, \mathfrak{a}'$ in $\mathcal{S}_{1,f}$. By virtue of Lemma 36, this means $\mathfrak{a} \sim_1 \mathfrak{a}'$.

Because of $\mathcal{A} \models \chi_3$, there are two elements $\widetilde{\mathfrak{a}}, \widetilde{\mathfrak{a}}' \in \mathcal{I}_1$ for which (i) entails $\mathfrak{a} \sim_1 \widetilde{\mathfrak{a}}$ and $\mathfrak{a}' \sim_1 \widetilde{\mathfrak{a}}'$. By symmetry and transitivity of $\sim_1$, we have $\widetilde{\mathfrak{a}} \sim_1 \widetilde{\mathfrak{a}}'$. Hence, we observe

$$\mathcal{A}, [\widetilde{j} \mapsto \widetilde{\mathfrak{a}}, \widetilde{j}' \mapsto \widetilde{\mathfrak{a}}'] \models \forall \widetilde{i}.\, L(\underline{\ell-1}, \widetilde{i}) \to \big(J(\underline{\ell}, \widetilde{j}, \widetilde{i}, 0) \leftrightarrow J(\underline{\ell}, \widetilde{j}', \widetilde{i}, 0)\big) \,.$$

Consequently, $\mathcal{A} \models \chi_3$ leads to $\mathfrak{a} = \mathfrak{a}'$.

Ad (vi). As we have seen in the proof of Lemma 38, we can define a bijective mapping μ that maps the functions in $\mathcal{F}_0$ to bit strings of length μ that either have a 0 as most significant bit or correspond to $0 \ldots 01$ (where the most significant bit is the rightmost one).

Let $p := 2^{\mu-1} + 1$. By virtue of Lemma 38, we can uniquely construct a chain

$$f_1 \sqsubset_0 f_2 \sqsubset_0 \ldots \sqsubset_0 f_p$$

comprising all functions in $\mathcal{F}_0$.

Properties (iv) and (v) together yield that $\mathcal{I}_1 = \{\mathfrak{a}_1, \ldots, \mathfrak{a}_p\}$ where $\mathfrak{a}_k \in \mathcal{S}_{1,f_k}$ for every $k = 1, \ldots, p$. Lemma 35 says that for any $\mathfrak{a}_k, \mathfrak{a}_{k'}$ with $\mathfrak{a}_k \prec_1 \mathfrak{a}_{k'}$ we also observe $f_k \sqsubset_0 f_{k'}$. By definition of $\sqsubset_0$ and the fact that all the $f_k, f_{k'}$ are distinct, it clearly holds that $f_k \sqsubset_0 f_{k'}$ implies $k' = k+1$ and $1 \leq k < k' \leq p$. Hence, $\mathfrak{a}_k \prec_1 \mathfrak{a}_{k'}$ can only hold if $k' = k+1$ and $1 \leq k < k' \leq p$.

Consider any element $\mathfrak{a} \in \mathcal{I}_1$ for which $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto c_\mu^{\mathcal{A}}] \models J(\underline{1}, j, i, 0)$. By $\mathcal{A} \models \psi_{16}$, we know that there are elements $\widetilde{\mathfrak{a}}, \widetilde{\mathfrak{a}}' \in \mathcal{I}_1$ such that $\widetilde{\mathfrak{a}} \prec_1 \widetilde{\mathfrak{a}}'$ and $\mathcal{A}, [j \mapsto \mathfrak{a}, \widetilde{j} \mapsto \widetilde{\mathfrak{a}}] \models \mathrm{eq}^1_{j,\widetilde{j}}$. By (i), the latter translates to $\mathfrak{a} \sim_1 \widetilde{\mathfrak{a}}$.

Let $g, \widetilde{g}$ be functions such that $\mathfrak{a} \in \mathcal{S}_{1,g}$ and $\widetilde{\mathfrak{a}} \in \mathcal{S}_{1,\widetilde{g}}$. Such functions exist by virtue of (iv). However, $\mathfrak{a} \sim_1 \widetilde{\mathfrak{a}}$ entails $g = \widetilde{g}$, by Lemma 37. But then (v) leads to $\mathfrak{a} = \widetilde{\mathfrak{a}}$. Consequently, we have $\mathfrak{a} \prec_1 \widetilde{\mathfrak{a}}'$. By (v), this means that all but one element in $\mathcal{I}_1$ must have a successor in $\mathcal{I}_1$ and all but one elements in $\mathcal{I}_1$ are successor in $\mathcal{I}_1$. Hence, we obtain the chain $\mathfrak{a}_1 \prec_1 \mathfrak{a}_2 \prec_1 \ldots \prec_1 \mathfrak{a}_{p-1} \prec_1 \mathfrak{a}_p$ where $\mathfrak{a}_1 = d_1^{\mathcal{A}}$ and $\mathfrak{a}_p = e_1^{\mathcal{A}}$.

Inductive case $\ell > 1$.

Ad (i). Due to the assumption $\mathcal{A}, [j \mapsto \mathfrak{a}, \widetilde{j} \mapsto \widetilde{\mathfrak{a}}] \models \mathrm{eq}^\ell_{j,\widetilde{j}}$ with $\ell > 1$, the construction of $\mathrm{eq}^\ell_{j,\widetilde{j}}$ entails

$$\begin{aligned} \mathcal{A}, [j \mapsto \mathfrak{a}, \widetilde{j} \mapsto \widetilde{\mathfrak{a}}] \models\ \forall i.\, L(\underline{\ell-1}, i) \ \to\ & \exists \widetilde{i}.\, L(\underline{\ell-1}, \widetilde{i}) \wedge \mathrm{eq}^{\ell-1}_{i,\widetilde{i}} \\ & \wedge \Big(\big(J(\underline{\ell}, j, i, 0) \leftrightarrow J(\underline{\ell}, \widetilde{j}, \widetilde{i}, 0)\big) \\ & \quad \wedge \big(J(\underline{\ell}, j, i, 1) \leftrightarrow J(\underline{\ell}, \widetilde{j}, \widetilde{i}, 1)\big)\Big) \,. \end{aligned}$$

By inductive application of (i), $\mathcal{A}, [i{\mapsto}\mathfrak{b}, \widetilde{i}{\mapsto}\widetilde{\mathfrak{b}}] \models \mathrm{eq}^{\ell-1}_{i,\widetilde{i}}$ entails $\mathfrak{b} \sim_{\ell-1} \widetilde{\mathfrak{b}}$. Inductive application of (iv) implies that $\mathfrak{b} \in \mathcal{S}_{\ell-1,f}$ for some $f \in \mathcal{F}_{\ell-1}$. By Lemma 37 together with $\mathfrak{b} \sim_{\ell-1} \widetilde{\mathfrak{b}}$, we conclude $\widetilde{\mathfrak{b}} \in \mathcal{S}_{\ell-1,f}$. Now, inductive application of (v) leads to $\mathfrak{b} = \widetilde{\mathfrak{b}}$. This means we in fact have

$$\mathcal{A}, [j{\mapsto}\mathfrak{a}, \widetilde{j}{\mapsto}\widetilde{\mathfrak{a}}] \models \forall i.\, L(\underline{\ell-1}, i) \rightarrow \Big(\big(J(\underline{\ell}, j, i, 0) \leftrightarrow J(\underline{\ell}, \widetilde{j}, \widetilde{i}, 0)\big) \wedge \big(J(\underline{\ell}, j, i, 1) \leftrightarrow J(\underline{\ell}, \widetilde{j}, \widetilde{i}, 1)\big)\Big) .$$

In other words, $\mathfrak{a} \sim_\ell \widetilde{\mathfrak{a}}$.

Ad (ii). By definition of $\sim_\ell$, $\mathfrak{a} \sim_\ell \widetilde{\mathfrak{a}}$ entails that for every $\mathfrak{b} \in \mathcal{I}_{\ell-1}$ we have

- $\mathcal{A}, [j{\mapsto}\mathfrak{a}, i{\mapsto}\mathfrak{b}] \models J(\underline{1}, j, i, 0)$ if and only if $\mathcal{A}, [\widetilde{j}{\mapsto}\widetilde{\mathfrak{a}}, i{\mapsto}\mathfrak{b}] \models J(\underline{1}, \widetilde{j}, i, 0)$ and
- $\mathcal{A}, [j{\mapsto}\mathfrak{a}, i{\mapsto}\mathfrak{b}] \models J(\underline{1}, j, i, 1)$ if and only if $\mathcal{A}, [\widetilde{j}{\mapsto}\widetilde{\mathfrak{a}}, i{\mapsto}\mathfrak{b}] \models J(\underline{1}, \widetilde{j}, i, 1)$.

By (i) and the fact that $\sim_{\ell-1}$ is an equivalence relation and thus $\mathfrak{b} \sim_{\ell-1} \mathfrak{b}$, we conclude $\mathcal{A}, [i{\mapsto}\mathfrak{b}, \widetilde{i}{\mapsto}\mathfrak{b}] \models \mathrm{eq}^{\ell-1}_{i,\widetilde{i}}$ for every $\mathfrak{b} \in \mathcal{I}_{\ell-1}$. Consequently, $\mathcal{A}, [j{\mapsto}\mathfrak{a}, \widetilde{j}{\mapsto}\widetilde{\mathfrak{a}}] \models \mathrm{eq}^{\ell}_{j,\widetilde{j}}$.

Ad (iii). By inductive application of (vi), we know that there is a unique chain $\mathfrak{b}_1 \prec_{\ell-1} \ldots \prec_{\ell-1} \mathfrak{b}_p$ comprising all elements in $\mathcal{I}_{\ell-1}$. Moreover, we observe $\mathcal{A}, [i{\mapsto}\mathfrak{b}_1] \models \mathrm{MinIdx}(\underline{\ell-1}, i)$ and $\mathcal{A}, [i{\mapsto}\mathfrak{b}_p] \models \mathrm{MaxIdx}(\underline{\ell-1}, i)$. Hence, by Lemma 38, we can arrange all mappings in $\mathcal{F}_{\ell-1}$ into a sequence $f_1 \sqsubset_{\ell-1} \ldots \sqsubset_{\ell-1} f_{p'}$ where $p' = 2^{p-1}+1$.

Clearly, f_1 maps every element $\mathfrak{b} \in \mathcal{I}_{\ell-1}$ to $f_1(\mathfrak{b}) = 0$, and $f_{p'}$ maps every element $\mathfrak{b}_k$ with $k < p$ to $f_{p'}(\mathfrak{b}_k) = 0$ and $\mathfrak{b}_p$ to $f_{p'}(\mathfrak{b}_p) = 1$. By $\mathcal{A} \models \psi_3 \wedge \psi_{10}$, we know that $\mathcal{A}, [j{\mapsto}d^{\mathcal{A}}_\ell, i{\mapsto}\mathfrak{b}] \models J(\underline{\ell}, j, i, 0)$ for every $\mathfrak{b} \in \mathcal{I}_{\ell-1}$. Hence, $d^{\mathcal{A}}_\ell \in S_{\ell,f_1}$.

We next show that for every k, $1 \leq k < p'$, if S_{ℓ,f_k} is nonempty, then $S_{\ell,f_{k+1}}$ is nonempty. Let $\mathfrak{a}$ be an element of S_{ℓ,f_k}. Because of $k < p'$, we know that $\mathcal{A}, [j{\mapsto}\mathfrak{a}, i{\mapsto}\mathfrak{b}_p] \models J(\underline{\ell}, j, i, 0)$. By virtue of (i) and due to $\mathcal{A} \models \psi_{16}$ we conclude that there are elements $\widetilde{\mathfrak{a}}, \widetilde{\mathfrak{a}}' \in \mathcal{I}_\ell$ such that $\mathfrak{a} \sim_\ell \widetilde{\mathfrak{a}}$ and $\widetilde{\mathfrak{a}} \prec_\ell \widetilde{\mathfrak{a}}'$. Moreover, Lemma 37 leads to $\widetilde{\mathfrak{a}} \in S_{\ell,f_k}$. By Lemma 35, this results in $\widetilde{\mathfrak{a}}' \in \mathcal{S}_{\ell,f_{k+1}}$.

Ad (iv). As a consequence of $\mathcal{A} \models \psi_{13}$ there is a unique partial mapping $g_{\mathfrak{a}} : \mathcal{I}_{\ell-1} \to \{0,1\}$ for every $\mathfrak{a} \in \mathcal{I}_\ell$ such that for every $\mathfrak{b} \in \mathcal{I}_{\ell-1}$ we have $\mathcal{A}, [j{\mapsto}\mathfrak{a}, i{\mapsto}\mathfrak{b}] \models J(\underline{\ell}, j, i, g_{\mathfrak{a}}(\mathfrak{b}))$ if and only if $g_{\mathfrak{a}}$ is defined for $\mathfrak{b}$. Because of $\mathcal{A} \models \chi_2$, we know that $g_{\mathfrak{a}}$ must be total.

For every $\mathfrak{a} \in \mathcal{I}_\ell$ where $g_{\mathfrak{a}}(e^{\mathcal{A}}_{\ell-1}) = 0$ we have $g_{\mathfrak{a}} \in \mathcal{F}_{\ell-1}$ and thus also $\mathfrak{a} \in \mathcal{S}_{\ell,g_{\mathfrak{a}}}$.

Because of $\mathcal{A} \models \psi_5 \wedge \psi_{11} \wedge \psi_{12}$ we know that $\mathcal{A} \models \mathrm{MaxIdx}(\underline{\ell}, e_\ell)$ and that $e^{\mathcal{A}}_\ell$ is the only element in $\mathcal{I}_\ell$ for which $\mathcal{A}, [j{\mapsto}e^{\mathcal{A}}_\ell, i{\mapsto}e^{\mathcal{A}}_{\ell-1}] \models J(\underline{\ell}, j, i, 1)$.

It remains to show that for every $\mathfrak{b} \in \mathcal{I}_{\ell-1}$ with $\mathfrak{b} \neq e^{\mathcal{A}}_{\ell-1}$ we have $\mathcal{A}, [j{\mapsto}e^{\mathcal{A}}_\ell, i{\mapsto}\mathfrak{b}] \models J(\underline{\ell}, j, i, 0)$. But this is a consequence of (iii) and the fact that the total function f_* mapping all elements $\mathfrak{b}$ in $\mathcal{I}_{\ell-1}$ but $e^{\mathcal{A}}_{\ell-1}$ to $f_*(\mathfrak{b}) = 0$ belongs to $\mathcal{F}_{\ell-1}$, and thus $\mathcal{S}_{\ell,f_*}$ is nonempty. In particular, $\mathcal{S}_{\ell,f_*} = \{e^{\mathcal{A}}_\ell\}$.

Consequently, $\mathcal{I}_\ell$ cannot contain any elements that do not lie in $\bigcup_{f \in \mathcal{F}_{\ell-1}} \mathcal{S}_{\ell,f}$.

Ad (v). Consider any set $\mathcal{S}_{\ell,f}$. By virtue of (iii), $\mathcal{S}_{\ell,f}$ contains at least one element. Suppose we are given two elements $\mathfrak{a}, \mathfrak{a}'$ in $\mathcal{S}_{\ell,f}$. By virtue of Lemma 36, this means $\mathfrak{a} \sim_\ell \mathfrak{a}'$.

Because of $\mathcal{A} \models \chi_3$, there are two elements $\widetilde{\mathfrak{a}}, \widetilde{\mathfrak{a}}' \in \mathcal{I}_\ell$ for which (i) entails $\mathfrak{a} \sim_\ell \widetilde{\mathfrak{a}}$ and $\mathfrak{a}' \sim_\ell \widetilde{\mathfrak{a}}'$. By symmetry and transitivity of $\sim_\ell$, we have $\widetilde{\mathfrak{a}} \sim_\ell \widetilde{\mathfrak{a}}'$. Hence, we observe

$$\mathcal{A}, [\widetilde{j}{\mapsto}\widetilde{\mathfrak{a}}, \widetilde{j}'{\mapsto}\widetilde{\mathfrak{a}}'] \models \forall \widetilde{i}.\, L(\underline{\ell-1}, \widetilde{i}) \rightarrow \big(J(\underline{\ell}, \widetilde{j}, \widetilde{i}, 0) \leftrightarrow J(\underline{\ell}, \widetilde{j}', \widetilde{i}, 0)\big) .$$

Consequently, $\mathcal{A} \models \chi_3$ leads to $\mathfrak{a} = \mathfrak{a}'$.

Ad (vi). As we have seen in the proof of Lemma 38, we can define a bijective mapping μ that maps the functions in $\mathcal{F}_{\ell-1}$ to bit strings of length μ that either have a 0 as most significant bit or correspond to $0 \ldots 01$ (where the most significant bit is the rightmost one).

Let $p' := 2^{|\mathcal{I}_{\ell-1}|-1}+1$. By virtue of Lemma 38, we can uniquely construct a chain

$$f_1 \sqsubset_{\ell-1} f_2 \sqsubset_{\ell-1} \ldots \sqsubset_{\ell-1} f_{p'}$$

comprising all functions in $\mathcal{F}_{\ell-1}$.
Properties (iv) and (v) together yield that $\mathcal{I}_\ell = \{\mathfrak{a}_1, \ldots, \mathfrak{a}_{p'}\}$ where $\mathfrak{a}_k \in \mathcal{S}_{\ell,f_k}$ for every $k = 1, \ldots, p'$. Lemma 35 says that for any $\mathfrak{a}_k, \mathfrak{a}_{k'}$ with $\mathfrak{a}_k \prec_\ell \mathfrak{a}_{k'}$ we observe $f_k \sqsubset_{\ell-1} f_{k'}$. By definition of $\sqsubset_{\ell-1}$ and the fact that all the $f_k, f_{k'}$ are distinct, it clearly holds that $f_k \sqsubset_{\ell-1} f_{k'}$ implies $k' = k+1$ and $1 \leq k < k' \leq p'$. Hence, $\mathfrak{a}_k \prec_\ell \mathfrak{a}_{k'}$ can only hold if $k' = k+1$ and $1 \leq k < k' \leq p'$.
Consider any element $\mathfrak{a} \in \mathcal{I}_\ell$ for which $\mathcal{A}, [j \mapsto \mathfrak{a}, i \mapsto e^{\mathcal{A}}_{\ell-1}] \models J(\underline{\ell}, j, i, 0)$. By $\mathcal{A} \models \psi_{16}$, we know that there are elements $\widetilde{\mathfrak{a}}, \widetilde{\mathfrak{a}}' \in \mathcal{I}_\ell$ such that $\widetilde{\mathfrak{a}} \prec_\ell \widetilde{\mathfrak{a}}'$ and $\mathcal{A}, [j \mapsto \mathfrak{a}, \widetilde{j} \mapsto \widetilde{\mathfrak{a}}] \models \mathrm{eq}^{\ell}_{j,\widetilde{j}}$. By (i), the latter translates to $\mathfrak{a} \sim_\ell \widetilde{\mathfrak{a}}$.
Let $g, \widetilde{g}$ be functions such that $\mathfrak{a} \in \mathcal{S}_{\ell,g}$ and $\widetilde{\mathfrak{a}} \in \mathcal{S}_{\ell,\widetilde{g}}$. Such functions exist by virtue of (iv). However, $\mathfrak{a} \sim_\ell \widetilde{\mathfrak{a}}$ entails $g = \widetilde{g}$, by Lemma 37. Moreover, (v) leads to $\mathfrak{a} = \widetilde{\mathfrak{a}}$. Consequently, we have $\mathfrak{a} \prec_\ell \widetilde{\mathfrak{a}}'$. By (v), this means that all but one element in $\mathcal{I}_\ell$ must have a successor in $\mathcal{I}_\ell$ and all but one element in $\mathcal{I}_\ell$ are successor in $\mathcal{I}_\ell$. Hence, we obtain the chain $\mathfrak{a}_1 \prec_\ell \mathfrak{a}_2 \prec_\ell \ldots \prec_\ell \mathfrak{a}_{p'-1} \prec_\ell \mathfrak{a}_{p'}$ where $\mathfrak{a}_1 = d^{\mathcal{A}}_\ell$ and $\mathfrak{a}_{p'} = e^{\mathcal{A}}_\ell$. □

The following is a corollary of the previous lemma, and so is Lemma 24.

Corollary 41. For every $\ell = 1, \ldots, \kappa$ we have $|\mathcal{I}_\ell| = p$ where $p := 2^{|\mathcal{I}_{\ell-1}|-1}+1 = 2^{\uparrow \ell}(\mu-1)+1$. Moreover, there is a unique chain $\mathfrak{a}_1 \prec_\ell \ldots \prec_\ell \mathfrak{a}_p$ comprising all elements in $\mathcal{I}_\ell$, and $\mathcal{A}, [j \mapsto \mathfrak{a}_1] \models \mathrm{MinIdx}(\underline{\ell}, j)$ and $\mathcal{A}, [j' \mapsto \mathfrak{a}_p] \models \mathrm{MaxIdx}(\underline{\ell}, j')$.

F Details regarding Section 5.2

The following sentences encode a given domino problem $\mathfrak{D} := \langle \mathcal{D}, \mathcal{H}, \mathcal{V} \rangle$ plus an initial condition $\overline{D}$, which is a finite word over $\mathcal{D}$. We try to make as many sentences as possible Horn.

$$\eta_1 := \forall xyx'y'.\ H(x,y,x',y') \ \rightarrow\ L(\underline{\kappa},x) \wedge L(\underline{\kappa},y) \wedge L(\underline{\kappa},x') \wedge L(\underline{\kappa},y') \wedge y \approx y'$$

$$\eta_2 := \forall xyx'y'i.\ H(x,y,x',y') \wedge \mathrm{MaxIdx}(\underline{\kappa-1},i) \wedge J(\underline{\kappa},x,i,0) \ \rightarrow\ \mathrm{Succ}(\underline{\kappa},x,x')$$

$$\eta_3 := \forall xyi.\ L(\underline{\kappa},x) \wedge L(\underline{\kappa},y) \wedge \mathrm{MaxIdx}(\underline{\kappa-1},i) \wedge J(\underline{\kappa},x,i,0) \rightarrow\ \exists \widetilde{x}\,\widetilde{y}\,\widetilde{x}'.\ \mathrm{eq}^{\kappa}_{x,\widetilde{x}} \wedge \mathrm{eq}^{\kappa}_{y,\widetilde{y}} \wedge \Big(\bigwedge_{D \in \mathcal{D}} \underline{D}(x,y) \leftrightarrow \underline{D}(\widetilde{x},\widetilde{y}) \Big) \wedge H(\widetilde{x},\widetilde{y},\widetilde{x}',\widetilde{y})$$

$$\eta_4 := \forall xyx'.\ \mathrm{MaxIdx}(\underline{\kappa},x) \wedge \mathrm{MinIdx}(\underline{\kappa},x') \ \rightarrow\ H(x,y,x',y)$$

$$\eta_5 := \forall xyx'y'.\ H(x,y,x',y') \wedge \mathrm{MaxIdx}(\underline{\kappa},x) \ \rightarrow\ \mathrm{MinIdx}(\underline{\kappa},x')$$

$$\eta_6 := \forall xyx'y'.\ H(x,y,x',y') \wedge \mathrm{MinIdx}(\underline{\kappa},x') \ \rightarrow\ \mathrm{MaxIdx}(\underline{\kappa},x)$$

$$\eta_7 := \forall xyx'y'.\ V(x,y,x',y') \ \rightarrow\ L(\underline{\kappa},x) \wedge L(\underline{\kappa},y) \wedge L(\underline{\kappa},x') \wedge L(\underline{\kappa},y') \wedge x \approx x'$$

$$\eta_8 := \forall xyx'y'.\ V(x,y,x',y') \wedge \mathrm{MaxIdx}(\underline{\kappa-1},i) \wedge J(\underline{\kappa},y,i,0) \ \rightarrow\ \mathrm{Succ}(\underline{\kappa},y,y')$$

$$\eta_9 := \forall xyi.\ L(\underline{\kappa},x) \wedge L(\underline{\kappa},y) \wedge \mathrm{MaxIdx}(\underline{\kappa-1},i) \wedge J(\underline{\kappa},y,i,0) \rightarrow\ \exists \widetilde{x}\,\widetilde{y}\,\widetilde{y}'.\ \mathrm{eq}^{\kappa}_{x,\widetilde{x}} \wedge \mathrm{eq}^{\kappa}_{y,\widetilde{y}} \wedge \Big(\bigwedge_{D \in \mathcal{D}} \big(\underline{D}(x,y) \leftrightarrow \underline{D}(\widetilde{x},\widetilde{y})\big) \Big) \wedge V(\widetilde{x},\widetilde{y},\widetilde{x},\widetilde{y}')$$

$$\eta_{10} := \forall xyy'.\ \mathrm{MaxIdx}(\underline{\kappa},y) \wedge \mathrm{MinIdx}(\underline{\kappa},y') \ \rightarrow\ V(x,y,x,y')$$

$$\eta_{11} := \forall xx'yy'.\ V(x,y,x',y') \wedge \mathrm{MaxIdx}(\underline{\kappa},y) \ \rightarrow\ \mathrm{MinIdx}(\underline{\kappa},y')$$

$$\eta_{12} := \forall xx'yy'.\ V(x,y,x',y') \wedge \mathrm{MinIdx}(\underline{\kappa},y') \ \rightarrow\ \mathrm{MaxIdx}(\underline{\kappa},y)$$

$$\eta_{13} := \bigwedge_{D \in \mathcal{D}} \forall xy.\ \underline{D}(x,y) \ \rightarrow\ L(\underline{\kappa},x) \wedge L(\underline{\kappa},y)$$

$$\eta_{14} := \bigwedge_{D \in \mathcal{D}} \bigwedge_{D' \in \mathcal{D} \setminus \{D\}} \forall xy.\ \underline{D}(x,y) \ \rightarrow\ \neg \underline{D}'(x,y)$$

$$\eta_{15} := \forall xx'y.\, H(x,y,x',y) \;\rightarrow\; \bigvee_{\langle D,D'\rangle\in\mathcal{H}} \underline{D}(x,y)\wedge\underline{D}'(x',y)$$

$$\eta_{16} := \forall xyy'.\, V(x,y,x,y') \;\rightarrow\; \bigvee_{\langle D,D'\rangle\in\mathcal{V}} \underline{D}(x,y)\wedge\underline{D}'(x,y')$$

$$\eta_{17} := \forall z.\, \mathrm{MinIdx}(\underline{\kappa},z) \;\rightarrow\; f_1 \approx z \wedge \bigwedge_{i=1}^{n-1} H(f_i,z,f_{i+1},z)$$

$$\eta_{18} := \forall z.\, \mathrm{MinIdx}(\underline{\kappa},z) \;\rightarrow\; \bigwedge_{i=1}^{n} \underline{D}_i(f_i,z)$$

The sentences η_{17} and η_{18} express the initial condition $\overline{D}$ where the lower left domino tiles are predefined to be the sequence $\overline{D} = D_1 \dots D_n$.

Regarding the length of the sentences $\eta_1,\dots,\eta_{18}$, we observe the following:

- $\mathrm{len}(\eta_1),\mathrm{len}(\eta_2),\mathrm{len}(\eta_4),\dots,\mathrm{len}(\eta_8),\mathrm{len}(\eta_{10}),\dots,\mathrm{len}(\eta_{12}) \in \mathcal{O}(\log\kappa)$
- $\mathrm{len}(\eta_3),\mathrm{len}(\eta_9) \in \mathcal{O}\big(\kappa\log\kappa + \mu(\log\kappa+\log\mu) + |\mathcal{D}|\log|\mathcal{D}|\big)$
- $\mathrm{len}(\eta_{13}) \in \mathcal{O}\big(|\mathcal{D}|(\log|\mathcal{D}|+\log\kappa)\big)$
- $\mathrm{len}(\eta_{14}) \in \mathcal{O}\big(|\mathcal{D}|^2\log|\mathcal{D}|\big)$
- $\mathrm{len}(\eta_{15}) \in \mathcal{O}\big(|\mathcal{H}|\log|\mathcal{D}|\big) = \mathcal{O}\big(|\mathcal{D}|^2\log|\mathcal{D}|\big)$
- $\mathrm{len}(\eta_{16}) \in \mathcal{O}\big(|\mathcal{H}|\log|\mathcal{D}|\big) = \mathcal{O}\big(|\mathcal{D}|^2\log|\mathcal{D}|\big)$
- $\mathrm{len}(\eta_{17}) \in \mathcal{O}\big(\log\kappa + n\log n\big)$
- $\mathrm{len}(\eta_{18}) \in \mathcal{O}\big(\log\kappa + n(\log n + \log|\mathcal{D}|)\big)$

In total, the length of $\eta_1\wedge\dots\wedge\eta_{18}$ lies in $\mathcal{O}\big(\widehat{n}\log\widehat{n}\big)$, where $\widehat{n} := \max\{\kappa,\mu,n,|\mathcal{D}|^2\}$.

For the remainder of this subsection we consider a model $\mathcal{A}$ of the sentence $\psi_1\wedge\dots\wedge\psi_{16}\wedge\chi_1\wedge\chi_2\wedge\chi_3\wedge\eta_1\wedge\dots\wedge\eta_{18}$. Moreover, we take over the notation from Definition 23.

Definition 42. We define the following relations:

- $\prec^H \subseteq \mathcal{I}_\kappa^2\times\mathcal{I}_\kappa^2$ such that $\langle\mathfrak{a},\mathfrak{b}\rangle \prec^H \langle\mathfrak{a}',\mathfrak{b}'\rangle$ holds if and only if $\mathcal{A},[x{\mapsto}\mathfrak{a},y{\mapsto}\mathfrak{b},x'{\mapsto}\mathfrak{a}',y'{\mapsto}\mathfrak{b}'] \models H(x,y,x',y')$.
- $\prec^V \subseteq \mathcal{I}_\kappa^2\times\mathcal{I}_\kappa^2$ such that $\langle\mathfrak{a},\mathfrak{b}\rangle \prec^V \langle\mathfrak{a}',\mathfrak{b}'\rangle$ holds if and only if $\mathcal{A},[x{\mapsto}\mathfrak{a},y{\mapsto}\mathfrak{b},x'{\mapsto}\mathfrak{a}',y'{\mapsto}\mathfrak{b}'] \models V(x,y,x',y')$.

Lemma 43. For all pairs $\langle\mathfrak{a},\mathfrak{b}\rangle,\langle\mathfrak{a}',\mathfrak{b}'\rangle \in \mathcal{I}_\kappa^2$ we observe the following properties.

(i) $\langle\mathfrak{a},\mathfrak{b}\rangle \prec^H \langle\mathfrak{a}',\mathfrak{b}'\rangle$ entails that $\mathfrak{b}=\mathfrak{b}'$ and that either $\mathfrak{a}\prec_\kappa\mathfrak{a}'$ or $\mathfrak{a}=e_\kappa^{\mathcal{A}}$ and $\mathfrak{a}'=d_\kappa^{\mathcal{A}}$.

(ii) $\mathfrak{a}=e_\kappa^{\mathcal{A}}$ and $\mathfrak{a}'=d_\kappa^{\mathcal{A}}$ entails $\langle\mathfrak{a},\mathfrak{b}\rangle \prec^H \langle\mathfrak{a}',\mathfrak{b}\rangle$ for every $\mathfrak{b}\in\mathcal{I}_\kappa$.

(iii) $\mathfrak{a}\prec_\kappa\mathfrak{a}'$ implies $\langle\mathfrak{a},\mathfrak{b}\rangle \prec^H \langle\mathfrak{a}',\mathfrak{b}\rangle$ for every $\mathfrak{b}\in\mathcal{I}_\kappa$.

(iv) $\langle\mathfrak{a},\mathfrak{b}\rangle \prec^V \langle\mathfrak{a}',\mathfrak{b}'\rangle$ entails that $\mathfrak{a}=\mathfrak{a}'$ and that either $\mathfrak{b}\prec_\kappa\mathfrak{b}'$ or $\mathfrak{b}=e_\kappa^{\mathcal{A}}$ and $\mathfrak{b}'=d_\kappa^{\mathcal{A}}$.

(v) $\mathfrak{b}=e_\kappa^{\mathcal{A}}$ and $\mathfrak{b}'=d_\kappa^{\mathcal{A}}$ entails $\langle\mathfrak{a},\mathfrak{b}\rangle \prec^V \langle\mathfrak{a},\mathfrak{b}'\rangle$ for every $\mathfrak{a}\in\mathcal{I}_\kappa$.

(vi) $\mathfrak{b}\prec_\kappa\mathfrak{b}'$ implies $\langle\mathfrak{a},\mathfrak{b}\rangle \prec^V \langle\mathfrak{a},\mathfrak{b}'\rangle$ for every $\mathfrak{a}\in\mathcal{I}_\kappa$.

Proof. Property (i) follows by $\mathcal{A} \models \eta_1 \wedge \eta_2 \wedge \eta_5 \wedge \eta_6$. Property (ii) follows by $\mathcal{A} \models \eta_4$.

In order to show Property (iii), we have to argue a bit more. First of all, we conclude $\mathfrak{a} \neq e_\kappa^{\mathcal{A}}$, by $\mathcal{A} \models \psi_2 \wedge \psi_4$. Because of $\mathcal{A} \models \eta_3$ and Lemma 40(i), for all $\mathfrak{a}, \mathfrak{b}$ with $\mathfrak{a} \neq e_\kappa^{\mathcal{A}}$ there must exist $\widetilde{\mathfrak{a}}, \widetilde{\mathfrak{b}}, \widetilde{\mathfrak{a}}' \in \mathcal{I}_\kappa$ such that $\mathfrak{a} \sim_\kappa \widetilde{\mathfrak{a}}$, $\mathfrak{b} \sim_\kappa \widetilde{\mathfrak{b}}$, and $\langle \widetilde{\mathfrak{a}}, \widetilde{\mathfrak{b}} \rangle \prec^H \langle \widetilde{\mathfrak{a}}', \widetilde{\mathfrak{b}} \rangle$. Due to Lemma 37 in combination with Lemma 40(v), we get $\mathfrak{a} = \widetilde{\mathfrak{a}}$ and $\mathfrak{b} = \widetilde{\mathfrak{b}}$. Thus, we have $\langle \mathfrak{a}, \mathfrak{b} \rangle \prec^H \langle \widetilde{\mathfrak{a}}', \mathfrak{b} \rangle$. Moreover, $\mathfrak{a} \neq e_\kappa^{\mathcal{A}}$ together with (i) leads to $\mathfrak{a} \prec_\kappa \widetilde{\mathfrak{a}}'$. Since we assumed $\mathfrak{a} \prec_\kappa \mathfrak{a}'$, Lemma 40(vi) says that $\mathfrak{a}'$ is the only element satisfying $\mathfrak{a} \prec_\kappa \mathfrak{a}'$, i.e. $\widetilde{\mathfrak{a}}' = \mathfrak{a}'$. Consequently, it in fact holds $\langle \mathfrak{a}, \mathfrak{b} \rangle \prec^H \langle \mathfrak{a}', \mathfrak{b} \rangle$.

Properties (iv) to (vi) can be proved analogously to the first three properties using $\mathcal{A} \models \eta_7 \wedge \ldots \wedge \eta_{12}$. □

Lemma 44. Let $r := 2^{\uparrow\kappa}(\mu - 1) + 1$. There is a bijective mapping $\rho : \mathbb{Z}_r^2 \to \mathcal{I}_\kappa^2$ such that $\rho(0,0) = \langle d_\kappa^{\mathcal{A}}, d_\kappa^{\mathcal{A}} \rangle$ and for every pair $\langle s, t \rangle \in \mathbb{Z}_r^2$ we have

- $\rho(s,t) \prec^H \rho(s+1,t)$ and
- $\rho(s,t) \prec^V \rho(s,t+1)$.

where $+$ stands for addition modulo r.

Proof. By Corollary 41 we know that there is a unique chain $\mathfrak{a}_1 \prec_\kappa \ldots \prec_\kappa \mathfrak{a}_r$ comprising all elements in $\mathcal{I}_\kappa$. Notice that $\mathfrak{a}_k, \mathfrak{a}_{k'}$ with $k \neq k'$ are distinct. We define ρ so that $\rho(s,t) := \langle \mathfrak{a}_{s+1}, \mathfrak{a}_{t+1} \rangle$ for all $s, t \in \mathbb{Z}_r = \{0, \ldots r-1\}$.

Obviously, ρ is bijective. Since $\mathfrak{a}_1$ is the only element in $\mathcal{I}_\kappa$ for which there is no $\mathfrak{b}$ in the above chain with $\mathfrak{b} \prec_\kappa \mathfrak{a}_1$, $\mathcal{A} \models \psi_2 \wedge \psi_3$ enforces $\mathfrak{a}_1 = d_\kappa^{\mathcal{A}}$. Hence, $\rho(0,0) = \langle d_\kappa^{\mathcal{A}}, d_\kappa^{\mathcal{A}} \rangle$.

Since $\mathfrak{a}_r$ is the only element in $\mathcal{I}_\kappa$ for which there is no $\mathfrak{b}'$ in the above chain with $\mathfrak{a}_r \prec_\kappa \mathfrak{b}'$, $\mathcal{A} \models \psi_4 \wedge \psi_5$ enforces $\mathfrak{a}_r = e_\kappa^{\mathcal{A}}$. Hence, Lemma 43(ii) entails $\rho(r-1,t) \prec^H \rho(0,t)$ for every $t \in \mathbb{Z}_r$. Moreover, the existence of the above chain together with Lemma 43(iii) leads to $\rho(s,t) \prec^H \rho(s+1,t)$ for every $s \in \mathbb{Z}_r \setminus \{r-1\}$ and every $t \in \mathbb{Z}_r$. Consequently, we observe $\rho(s,t) \prec^H \rho(s+1,t)$—modulo r—for every pair $\langle s, t \rangle \in \mathbb{Z}_r^2$.

By similar arguments, we may infer $\rho(s,t) \prec^V \rho(s,t+1)$ for every pair $\langle s, t \rangle \in \mathbb{Z}_r^2$. □

Lemma 45. Assume that $\mathcal{D}$, $\mathcal{H}$, and $\mathcal{V}$ are nonempty. For all pairs $\langle \mathfrak{a}, \mathfrak{b} \rangle \in \mathcal{I}_\kappa$ we have $\mathcal{A}, [x \mapsto \mathfrak{a}, y \mapsto \mathfrak{b}] \models \underline{D}(x,y)$ for exactly one $D \in \mathcal{D}$.

Proof. Due to $\mathcal{A} \models \eta_{15} \wedge \eta_{16}$, we observe the following properties for all pairs $\langle \mathfrak{a}, \mathfrak{b} \rangle, \langle \mathfrak{a}', \mathfrak{b}' \rangle \in \mathcal{I}_\kappa^2$:

- $\langle \mathfrak{a}, \mathfrak{b} \rangle \prec^H \langle \mathfrak{a}', \mathfrak{b}' \rangle$ implies that there are $D, D' \in \mathcal{D}$ such that $\langle D, D' \rangle \in \mathcal{H}$ and $\mathcal{A}, [x \mapsto \mathfrak{a}, y \mapsto \mathfrak{b}] \models \underline{D}(x,y)$ and $\mathcal{A}, [x \mapsto \mathfrak{a}', y \mapsto \mathfrak{b}'] \models \underline{D'}(x,y)$.
- $\langle \mathfrak{a}, \mathfrak{b} \rangle \prec^V \langle \mathfrak{a}', \mathfrak{b}' \rangle$ implies that there are $D, D' \in \mathcal{D}$ such that $\langle D, D' \rangle \in \mathcal{V}$ and $\mathcal{A}, [x \mapsto \mathfrak{a}, y \mapsto \mathfrak{b}] \models \underline{D}(x,y)$ and $\mathcal{A}, [x \mapsto \mathfrak{a}', y \mapsto \mathfrak{b}'] \models \underline{D'}(x,y)$.

By virtue of Lemma 44, we know that there is a bijection ρ such that for every pair $\langle \mathfrak{a}, \mathfrak{b} \rangle$ in the image of ρ there is another pair $\langle \mathfrak{a}', \mathfrak{b}' \rangle$ such that $\langle \mathfrak{a}, \mathfrak{b} \rangle \prec^H \langle \mathfrak{a}', \mathfrak{b}' \rangle$ or $\langle \mathfrak{a}, \mathfrak{b} \rangle \prec^V \langle \mathfrak{a}', \mathfrak{b}' \rangle$. Since the image of ρ is $\mathcal{I}_\kappa^2$, this means that there is at least one $D \in \mathcal{D}$ for every pair $\langle \mathfrak{a}, \mathfrak{b} \rangle \in \mathcal{I}_\kappa^2$ such that $\mathcal{A}, [x \mapsto \mathfrak{a}, y \mapsto \mathfrak{b}] \models \underline{D}(x,y)$.

Finally, because of $\mathcal{A} \models \eta_{14}$ we know that there is at most one $D \in \mathcal{D}$ for every pair $\langle \mathfrak{a}, \mathfrak{b} \rangle \in \mathcal{I}_\kappa^2$ such that $\mathcal{A}, [x \mapsto \mathfrak{a}, y \mapsto \mathfrak{b}] \models \underline{D}(x,y)$. □

Lemma 25. Let $r := 2^{\uparrow\kappa}(\mu - 1) + 1$ and assume that $\mathcal{D}$, $\mathcal{H}$, and $\mathcal{V}$ are nonempty. $\mathcal{A}$ induces a tiling τ of $\mathbb{Z}_r^2$ with initial condition $\overline{D} := D_1, \ldots, D_n$.

Proof. Let ρ be a bijection according to Lemma 44. We define the mapping τ such that $\tau(s,t) := D$ if and only if $\rho(s,t) = \langle \mathfrak{a}, \mathfrak{b} \rangle$ and $\mathcal{A}, [x \mapsto \mathfrak{a}, y \mapsto \mathfrak{b}] \models \underline{D}(x,y)$. By Lemma 45, we know that τ is well defined.

By $\mathcal{A} \models \eta_{17} \wedge \eta_{18}$, we know that $\mathcal{A}, [x \mapsto \mathfrak{a}, y \mapsto \mathfrak{b}] \models \underline{D_i}(x,y)$ for $\langle \mathfrak{a}, \mathfrak{b} \rangle = \rho(i,0)$ and $i = 0, \ldots, n-1$. Hence, τ satisfies the initial condition.

By definition of ρ and because of $\mathcal{A} \models \eta_{15} \wedge \eta_{16}$, we observe the following:

- For every pair $\langle s,t\rangle \in \mathbb{Z}_r^2$ there are pairs $\langle D, D'\rangle \in \mathcal{H}$, $\langle \mathfrak{a}, \mathfrak{b}\rangle = \rho(s,t)$, and $\langle \mathfrak{a}', \mathfrak{b}'\rangle = \rho(s+1,t)$ such that
  - $\mathcal{A}, [x{\mapsto}\mathfrak{a}, y{\mapsto}\mathfrak{b}] \models \underline{D}(x,y)$ and
  - $\mathcal{A}, [x{\mapsto}\mathfrak{a}', y{\mapsto}\mathfrak{b}'] \models \underline{D'}(x,y)$.
- For every pair $\langle s,t\rangle \in \mathbb{Z}_r^2$ there are pairs $\langle D, D'\rangle \in \mathcal{V}$, $\langle \mathfrak{a}, \mathfrak{b}\rangle = \rho(s,t)$, and $\langle \mathfrak{a}', \mathfrak{b}'\rangle = \rho(s,t+1)$ such that
  - $\mathcal{A}, [x{\mapsto}\mathfrak{a}, y{\mapsto}\mathfrak{b}] \models \underline{D}(x,y)$ and
  - $\mathcal{A}, [x{\mapsto}\mathfrak{a}', y{\mapsto}\mathfrak{b}'] \models \underline{D'}(x,y)$.

Consequently, the mapping τ constitutes a proper tiling of $\mathbb{Z}_r^2$. □

G Details regarding Section 5.3

The following sentences comprise the axioms of reflexivity, symmetry, transitivity, and several axioms of congruence ($\psi'_1, \ldots, \psi'_{12}, \eta'_1, \ldots, \eta'_{10}$) for the binary predicate E that is introduced to replace equality.

$$\psi'_1 := \forall j.\ E(j,j)$$
$$\psi'_2 := \forall jj'.\ E(j,j') \ \rightarrow\ E(j',j)$$
$$\psi'_3 := \forall jj'j''.\ E(j,j') \wedge E(j',j'') \ \rightarrow\ E(j,j'')$$
$$\psi'_4 := \forall jj'.\ \bigwedge_{\ell=0}^{\kappa} E(j,j) \wedge L(\underline{\ell},j) \ \rightarrow\ L(\underline{\ell},j')$$
$$\psi'_5 := \forall jj'.\ \bigwedge_{\ell=0}^{\kappa} E(j,j) \wedge \mathrm{MinIdx}(\underline{\ell},j) \ \rightarrow\ \mathrm{MinIdx}(\underline{\ell},j')$$
$$\psi'_6 := \forall jj'.\ \bigwedge_{\ell=0}^{\kappa} E(j,j) \wedge \mathrm{MaxIdx}(\underline{\ell},j) \ \rightarrow\ \mathrm{MaxIdx}(\underline{\ell},j')$$
$$\psi'_7 := \forall jj'i.\ \bigwedge_{\ell=0}^{\kappa} E(j,j) \wedge \mathrm{Succ}(\underline{\ell},j,i) \ \rightarrow\ \mathrm{Succ}(\underline{\ell},j',i)$$
$$\psi'_8 := \forall jii'.\ \bigwedge_{\ell=0}^{\kappa} E(j,j) \wedge \mathrm{Succ}(\underline{\ell},j,i) \ \rightarrow\ \mathrm{Succ}(\underline{\ell},j,i')$$
$$\psi'_9 := \forall jj'i.\ \bigwedge_{\ell=0}^{\kappa} \big(E(j,j) \wedge J(\underline{\ell},j,i,0) \ \rightarrow\ J(\underline{\ell},j',i,0)\big) \wedge \big(E(j,j) \wedge J(\underline{\ell},j,i,1) \ \rightarrow\ J(\underline{\ell},j',i,1)\big)$$
$$\psi'_{10} := \forall jii'.\ \bigwedge_{\ell=0}^{\kappa} \big(E(j,j) \wedge J(\underline{\ell},j,i,0) \ \rightarrow\ J(\underline{\ell},j,i',0)\big) \wedge \big(E(j,j) \wedge J(\underline{\ell},j,i,1) \ \rightarrow\ J(\underline{\ell},j,i',1)\big)$$
$$\psi'_{11} := \forall jj'i.\ \bigwedge_{\ell=0}^{\kappa} \big(E(j,j) \wedge J^*(\underline{\ell},j,i,0) \ \rightarrow\ J^*(\underline{\ell},j',i,0)\big) \wedge \big(E(j,j) \wedge J^*(\underline{\ell},j,i,1) \ \rightarrow\ J^*(\underline{\ell},j',i,1)\big)$$
$$\psi'_{12} := \forall jii'.\ \bigwedge_{\ell=0}^{\kappa} \big(E(j,j) \wedge J^*(\underline{\ell},j,i,0) \ \rightarrow\ J^*(\underline{\ell},j,i',0)\big) \wedge \big(E(j,j) \wedge J^*(\underline{\ell},j,i,1) \ \rightarrow\ J^*(\underline{\ell},j,i',1)\big)$$

$$
\begin{aligned}
\eta'_1 &:= \forall xyuvx'.\, E(x,x') \wedge H(x,y,u,v) \;\rightarrow\; H(x',y,u,v)\\
\eta'_2 &:= \forall xyuvy'.\, E(y,y') \wedge H(x,y,u,v) \;\rightarrow\; H(x,y',u,v)\\
\eta'_3 &:= \forall xyuvu'.\, E(u,u') \wedge H(x,y,u,v) \;\rightarrow\; H(x,y,u',v)\\
\eta'_4 &:= \forall xyuvv'.\, E(v,v') \wedge H(x,y,u,v) \;\rightarrow\; H(x,y,u,v')\\
\eta'_5 &:= \forall xyuvx'.\, E(x,x') \wedge V(x,y,u,v) \;\rightarrow\; V(x',y,u,v)\\
\eta'_6 &:= \forall xyuvy'.\, E(y,y') \wedge V(x,y,u,v) \;\rightarrow\; V(x,y',u,v)\\
\eta'_7 &:= \forall xyuvu'.\, E(u,u') \wedge V(x,y,u,v) \;\rightarrow\; V(x,y,u',v)\\
\eta'_8 &:= \forall xyuvv'.\, E(v,v') \wedge V(x,y,u,v) \;\rightarrow\; V(x,y,u,v')\\
\eta'_9 &:= \forall xyx'.\, \bigwedge_{D\in\mathcal{D}} E(x,x') \wedge \underline{D}(x,y) \;\rightarrow\; \underline{D}(x',y)\\
\eta'_{10} &:= \forall xyy'.\, \bigwedge_{D\in\mathcal{D}} E(y,y') \wedge \underline{D}(x,y) \;\rightarrow\; \underline{D}(x,y')
\end{aligned}
$$

Regarding the length of the above sentences, we observe the following:

- $\mathrm{len}(\psi'_1), \mathrm{len}(\psi'_2), \mathrm{len}(\psi'_3) \in \mathcal{O}(1)$,
- $\mathrm{len}(\eta'_1), \ldots, \mathrm{len}(\eta'_8) \in \mathcal{O}(1)$,
- $\mathrm{len}(\psi'_4), \ldots, \mathrm{len}(\psi'_{12}) \in \mathcal{O}(\kappa \log \kappa)$,
- $\mathrm{len}(\eta'_9), \mathrm{len}(\eta'_{10}) \in \mathcal{O}\big(|\mathcal{D}| \log |\mathcal{D}|\big)$.

Lemma 26. Let $\varphi := \psi_1 \wedge \ldots \psi_{16} \wedge \chi_1 \wedge \chi_2 \wedge \chi_3 \wedge \eta_1 \wedge \ldots \wedge \eta_{18}$. Let $\widehat{\varphi}$ be the result of replacing every atom $s \approx t$ in φ with $E(s,t)$. Moreover, let $\psi' := \psi'_1 \wedge \ldots \wedge \psi'_{12}$ and $\eta' := \eta'_1 \wedge \ldots \wedge \eta'_{10}$.

(i) Every model $\mathcal{A} \models \varphi$ can be extended to a model $\mathcal{B} \models \widehat{\varphi} \wedge \psi' \wedge \eta'$ over the same domain.

(ii) From every model $\mathcal{B} \models \widehat{\varphi} \wedge \psi' \wedge \eta'$ we can construct a model $\mathcal{A} \models \varphi$.

Proof.

Ad (i). Let $\mathcal{A}$ be a model of φ. We define $\mathcal{B}$ exactly like $\mathcal{A}$ except for the interpretation of E, which we define to be $E^{\mathcal{B}} := \{\langle \mathfrak{a}, \mathfrak{a}\rangle \mid \mathfrak{a} \in \mathfrak{U}_{\mathcal{A}}\}$. Clearly, we have $\mathcal{B} \models \varphi$ and also $\mathcal{B} \models \widehat{\varphi}$. Moreover, it is easy to see that $\mathcal{B} \models \psi' \wedge \eta'$.

Ad (ii). Let $\mathcal{B}$ be a model of $\widehat{\varphi} \wedge \psi' \wedge \eta'$. For all domain elements $\mathfrak{a}, \mathfrak{b} \in \mathfrak{U}_{\mathfrak{B}}$ we write $\mathfrak{a} \sim_E \mathfrak{b}$ if and only if $\mathcal{B}, [j \mapsto \mathfrak{a}, j' \mapsto \mathfrak{b}] \models E(j,j')$. Due to $\mathcal{B} \models \psi'_1 \wedge \psi'_2 \wedge \psi'_3$, we know that $\sim_E$ constitutes an equivalence relation on $\mathfrak{U}_{\mathcal{B}}$. Let $\mathfrak{U}_{\mathcal{B}}/_{\sim_E}$ be the set of all equivalence classes induced by $\sim_E$.

Claim I: Consider two variable assignments β, β' such that for every variable j we have $\beta(j) \sim_E \beta'(j)$. For every atom A occurring in $\widehat{\varphi}$ we have $\mathcal{B}, \beta \models A$ if and only if $\mathcal{B}, \beta' \models A$.

Proof: We distinguish two cases: Let $A = P(\ldots)$, where P is one of the predicate symbols L, MinIdx, MaxIdx, Succ, J, J^*, H, V, or $\underline{D}$ with $D \in \mathcal{D}$. Then the claim follows due to $\mathcal{B} \models \psi' \wedge \eta'$ and the definition of $\sim_E$.

Let $A = E(j,i)$. Assume $\mathcal{B}, \beta \models E(j,i)$. Hence, we have $\beta'(j) \sim_E \beta(j) \sim_E \beta(i) \sim_E \beta'(i)$. By transitivity of $\sim_E$ it follows that $\beta'(j) \sim_E \beta'(i)$, and thus we have $\mathcal{B}, \beta' \models E(j,i)$.

Assume $\mathcal{B}, \beta' \models E(j,i)$. Hence, we have $\beta(j) \sim_E \beta'(j) \sim_E \beta'(i) \sim_E \beta(i)$. By transitivity of $\sim_E$ it follows that $\beta(j) \sim_E \beta(i)$, and thus we have $\mathcal{B}, \beta \models E(j,i)$. $\Diamond$

We define the structure $\mathcal{A}$ as follows:

- We define $\mathfrak{U}_{\mathcal{A}}$ such that it contains exactly one representative from every equivalence class in $\mathfrak{U}_{\mathcal{B}}/_{\sim_E}$.
- For every constant symbol c occurring in $\widehat{\varphi}$ we set $c^{\mathcal{A}} := \mathfrak{c}$, where $\mathfrak{c}$ is the representative in $\mathfrak{U}_{\mathcal{A}}$ such that it represents the $\sim_E$-class of $c^{\mathcal{B}}$, i.e. $[\mathfrak{c}]_{\sim_E} = [c^{\mathcal{B}}]_{\sim_E}$.
- For every predicate symbol P occurring in $\widehat{\varphi}$ we set $P^{\mathcal{A}} := P^{\mathcal{B}} \cap \mathfrak{U}_{\mathcal{A}}^m$, where m is the arity of P.

Claim II: $\mathcal{A} \models \widehat{\varphi}$.

Proof: By construction of $\mathfrak{U}_{\mathcal{A}}$, we can find for every variable assignment β mapping variables to elements in $\mathfrak{U}_{\mathcal{B}}$ another variable assignment β' mapping every variable to an element in $\mathfrak{U}_{\mathcal{A}}$ such that $\beta(j) \sim_E \beta'(j)$ holds for every variable j. Consequently, using Claim I, we can show that $\mathcal{A} \models \widehat{\varphi}$ holds by induction on the structure of $\widehat{\varphi}$. $\Diamond$

From the definition of $\mathcal{A}$ it follows that $E^{\mathcal{A}} = \{\langle \mathfrak{a}, \mathfrak{a} \rangle \mid \mathfrak{a} \in \mathfrak{U}_{\mathcal{A}}\}$. Hence, for all $\mathfrak{a}, \mathfrak{b} \in \mathfrak{U}_{\mathcal{A}}$ we have $\mathcal{A}, [j \mapsto \mathfrak{a}, j' \mapsto \mathfrak{b}] \models E(j, j')$ if and only if $\mathcal{A}, [j \mapsto \mathfrak{a}, j' \mapsto \mathfrak{b}] \models j \approx j'$. This observation together with Claim II entails $\mathcal{A} \models \varphi$. $\square$